

AD-A031 629

RAND CORP SANTA MONICA CALIF

F/G 5/1

A CRITICAL ASSESSMENT OF QUANTITATIVE METHODOLOGY AS A POLICY A--ETC(U)

AUG 74 R E STRAUCH

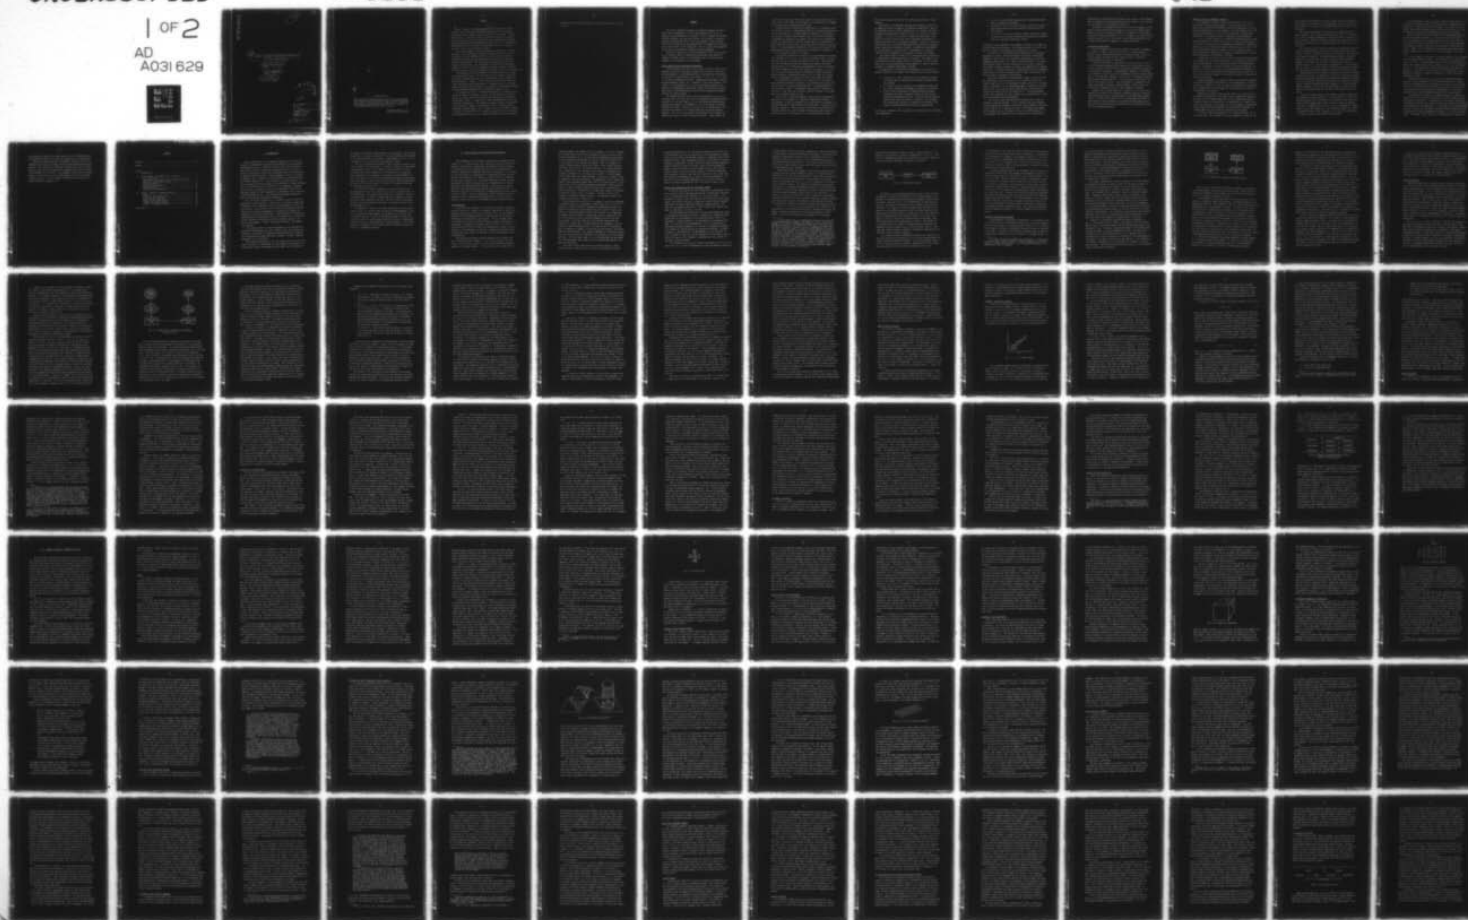
UNCLASSIFIED

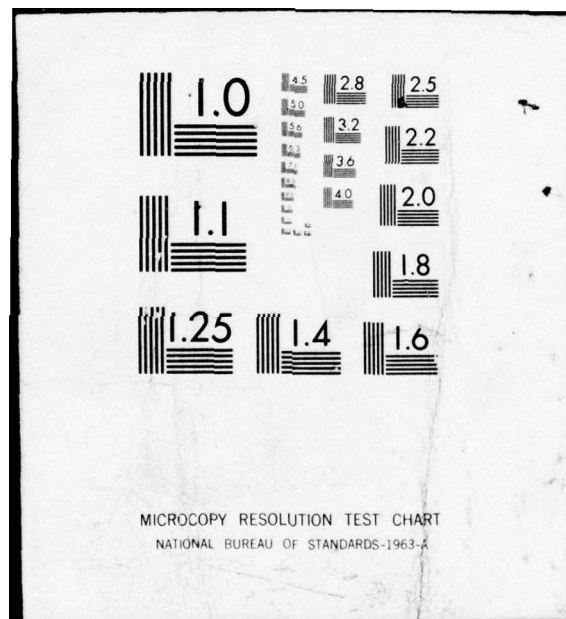
P-5282

•NL

1 OF 2

AD
A031 629





ADA031629

2

FL

6

A CRITICAL ASSESSMENT OF QUANTITATIVE METHODOLOGY
AS A POLICY ANALYSIS TOOL

10

Ralph E. /Strauch

11

August 1974

12

118p.

DDC
RECEIVED
NOV 5 1976
A

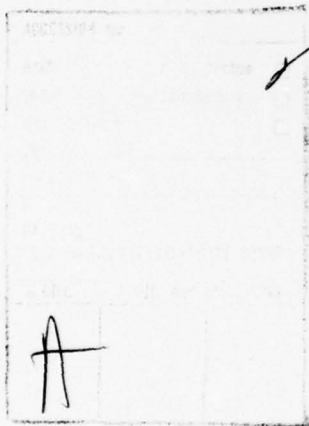
DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

14

P-5282

296600

LB



The Rand Paper Series

Papers are issued by The Rand Corporation as a service to its professional staff. Their purpose is to facilitate the exchange of ideas among those who share the author's research interests; Papers are not reports prepared in fulfillment of Rand's contracts or grants. Views expressed in a Paper are the author's own, and are not necessarily shared by Rand or its research sponsors.

The Rand Corporation ✓
Santa Monica, California 90406

PREFACE

Under a variety of names (such as ^{e.g.} operations analysis, systems analysis, and cost-benefit analysis) quantitative tools and methods have come to have a large and pervasive role in the analysis of public policy issues. The amount of effort devoted to quantitative analysis and the influence that this analysis has in the governmental planning and decision process make the question of what the role of quantitative methodology should be a question of legitimate concern for those who do, use, or rely on this type of analysis.

This paper presents a critical assessment of that role, focusing primarily on the limitations of quantitative methodology as a tool for the analysis of soft and "squishy" (i.e., without any well-defined mathematical formulation that unambiguously captures the substantive problem) problems and on the distortions that result when those limitations are neglected. This assessment is followed by a preliminary attempt to outline a theory of judgmental analysis intended to explore the difficulties more fully and suggest directions for their resolution.

The examples used as illustrations are drawn largely from military and defense applications. This bias should not be taken to imply that the issues addressed arise only, or even primarily, in the analysis of defense problems. Rather, it simply reflects the fact that my professional experience has been primarily as a defense analyst. The issues addressed arise equally strongly in the analysis of non-defense problems, and the assessment presented here is intended to apply to both.

This paper grows out of a longstanding concern on my part with the issues it attempts to address. As such, it has benefited from conversations and other forms of interaction with many people, both inside and outside of Rand. I regret that I cannot acknowledge all their contributions individually. For comments on earlier drafts of this paper, I am particularly indebted to Paul Berman, Garry D. Brewer, Herbert Goldhammer, Abraham Kaplan, John E. Koehler, Harold D. Lasswell, Edward S. Quade, and Charles Wolf, Jr. Needless to say, they do not all agree with everything said here, and I retain responsibility for the opinions

expressed and conclusions reached, as well as for any errors of fact or logic.

SUMMARY

The use of operations research and operations analysis in World War II was the beginning of a mushrooming expansion of the use of quantitative methodology in the analysis, first of defense and national security problems, and then more recently, of all manner of problems arising in governmental policy- and decisionmaking. Most advocates of the use of quantitative methodology view this expansion as salutary, and as an extension of the "objective" tools of science and mathematics to the arena of governmental planning and decisionmaking.

This report presents a critical assessment of that view, focusing on the inherent limits of quantitative methodology when applied to "squishy" problems, and on the unavoidable role of subjective human judgments in any analysis of such problems.

POLICY ANALYSIS AND QUANTITATIVE METHODOLOGY

In what follows, the term *policy analysis* is used to refer to the systematic analysis of questions arising from or faced by the governmental planning or decisionmaking process, conducted with the intention of affecting or contributing to that process. *Quantitative methodology* is the body of mathematical methods, computational techniques, supporting methodological theory, etc., available to the policy analyst. The utility of the methodology as a tool for application depends on the existence of appropriate similarities between the theoretical problems dealt with by the methodological theory and the substantive problems faced by the policy analyst.

The supporting theory can be classed as mathematics, per se, in the sense that it involves the study of mathematical systems and models as objects in their own right, i.e., as self-contained systems whose structure and behavior is determined by the set of axioms or premises defining the system. *Mathematical analysis* is the exploration of that structure and behavior as it follows logically from those defining premises. The *results* thus produced are grounded in the premises and connected to them by a chain of logical inference. These results are

"objective," in the sense that their validity can be determined on the basis of that chain alone, without appeal to the competence or judgment of the analyst who originally produced them.

The application of quantitative methodology to a substantive problem involves the use of a mathematical model as a simplified representation of that problem or of some phenomenon important to the problem. The conclusion reached should depend in part on the results produced by mathematical analysis of the model, and in part on the relationship between the problem and the model--on what parts of the problem the model represents, and how well, and on what parts it distorts, and how badly. This problem/model relationship falls outside the scope of the methodological theory, which deals with the analysis of problems *defined* by the models treated in the theory.

The simplest applications, relative to the problem/model relationship involved, are *rigorously quantifiable* problems in which the structure and logic of the problem and model are the same. Such problems arise in the analysis of physical systems obeying well understood physical laws, or in statistical experimentation in which the analyst introduces randomness to match that assumed in the model. This identity of structure between model and substantive problem means that results produced by mathematical analysis of the model can be given direct interpretation as substantive conclusions about the problem.

The conclusions produced thus are "objective," in the sense of being entirely grounded in the logic of the mathematical analysis and the empirical fit between problem and model. Like mathematical results, they are subject to independent verification or refutation on the basis of that logic and fit alone, without reference to the judgment of the analyst producing them.

This is not to imply that judgment plays no role in mathematical analysis or in the analysis of rigorously quantifiable problems. On the contrary, it plays an important role, particularly in original or creative analysis of complex and difficult problems. The nature of the problems, however, is such that, once results have been found or conclusions reached, the role of judgment in producing those results or conclusions can be "factored out" of the grounds supporting them.

This is not the case for many of the problems encountered in policy analysis.

While most of the problems occurring in the physical sciences or in engineering can be thought of as rigorously quantifiable, many of the problems arising in policy analysis are not. Rather, they span a spectrum from what might be called *reasonably quantifiable* to highly "squishy." *Reasonably quantifiable* problems are those with a high degree of natural structure that can be reasonably represented by a mathematical model, such as logistics problems. "Squishy" problems are those without any well-defined mathematical formulation that unambiguously captures the substantive problem. Formulations that appear to do so are squishy in the sense that the appearance of solidity is superficial and evaporates if the problem is leaned on or probed.

This report looks critically at the usefulness and the limitations of quantitative methodology across this spectrum. The emphasis is on the limitations at the squishy end of the spectrum. For this reason, some of the criticism may appear overly negative if taken to apply to analysis of reasonably quantifiable problems.

The application of quantitative methodology to the analysis of squishy problems can be thought of as having three levels, as follows:

1. A *substantive* level consisting of the substantive problem of interest and the conclusions the analysis provides about that problem.
2. A *mathematical* level consisting of a mathematical model and the results produced by mathematical analysis of that model.
3. A *formal* level consisting of a formal problem and formal conclusion that serve to mediate between and link the substantive and mathematical levels. These are seldom explicitly identified as distinct elements of analysis, but are always there implicitly, and must be separately articulated if one wishes to understand precisely what the analysis shows.

The process of analysis can be thought of as consisting of the following three components:

1. *Formulation* of the formal problem and the mathematical model used to represent the problem.
2. *Mathematical analysis* of and within the context of the model. This produces results that are logically valid within that that context.
3. *Interpretation* of those results as formal conclusions (within the context of the formal problem) and substantive conclusions about the underlying substantive problem.

In practice, these components blur together as aspects of a common interwoven intellectual process. They become separated, if at all, only when the results of that process are neaten up at the end.

The activities of formulation and interpretation are inherently judgmental in nature, and the conclusions reached are inextricably grounded in them. Hence, the conclusions themselves are inherently judgmental, particularly in the case of squishy problems. The methodology provides tools, but the usefulness of those tools in dealing with squishy problems depends on the skill and judgment of the analyst using them. To expect, as some do, the methodology itself to guarantee valid results is no more reasonable than to expect a particular type of paint brush to guarantee fine paintings.

The theory from which the methodology is drawn is concerned with the application of the methodology to the problems treated by the theory--those defined by the mathematical models with which the theory deals. Questions of how the methodology might be applied to problems outside that class (as most squishy problems are) fall outside the scope of the theory and are not treated by it. The methodological theory, then, provides no direct guidance for the application of judgment in formulation and interpretation.

This is not to say that quantitative methods have no value in analysis of squishy problems. On the contrary, they can frequently provide significant insights or solutions to such problems. When this occurs, however, it is usually due more to the insight and judgment of the analyst than to the power or rigor of the methodology per se. This suggests that standards by which applications of quantitative

methodology are judged should focus more on the quality of the judgments made in formulation and interpretation and less strictly on the technical correctness of the mathematical analysis.

One result of inadequate attention to the role of judgment is what might be called method-oriented analysis, in which the analyst simply pretends that the substantive problem is one amenable to his methodology, and proceeds accordingly. Thus, for example, the method-oriented regression analyst views the world as a multivariate Gaussian process in which every problem is amenable to regression analysis.

TWO FORMS OF OBJECTIVITY

One of the major reasons for the neglect of the role of judgment is a desire to emulate the "objectivity" of the natural sciences, without a clear understanding of what "objectivity" means in the context of squishy problems and what tradeoffs might be involved in achieving different forms of objectivity.

Scientific research is "objective" in the sense that it is *unprejudiced*--the analyst approaches his problem with an open mind, trying to "see things as they are" without preconceived prejudice forcing him to a particular conclusion. The knowledge he produces is "objective" in the sense of being *free-standing*--grounded in fact and logic independent of his subjective judgment, as in the case of rigorously quantifiable problems discussed earlier. These two forms of objectivity coincide in the natural sciences. This coincidence is a property of the class of problems the natural sciences address, however, and is not one that will necessarily occur in other classes of problems.

Indeed, there is good reason not to expect it to occur in policy analysis, where many of the problems encountered possess a political/social/behavioral content that man seems to best understand judgmentally. To strive for the *appearance* of free-standing objectivity in the analysis of such problems, as the method-oriented analyst often does, is itself a severe form of prejudice that can significantly distort the conclusions reached.

TOWARD A THEORY OF JUDGMENTAL ANALYSIS

This suggests a need for a better understanding of the role of careful and considered human judgment in analysis and of the use of quantitative methods and techniques as aids and supplements to, rather than as replacements for, that judgment. Some ideas around which such an understanding might be built are outlined below. They are now, admittedly, in a primitive stage of organization and development.

A *model*, broadly defined, is a simplified representation of a more complex phenomenon or system that we use to capture central features of, and help us understand, the phenomenon being modeled. Mathematical models are but one of the many types of models we use to do this. The use of mathematical models in analysis is, thus, a special case of a more general human cognitive use of models to understand the surrounding environment and the problems it presents.

A model can be used as a *surrogate* for the problem being addressed, in the sense that the structure and logic of the problem are taken as valid representations of the structure and logic of the substantive problem. The analyst, in effect, accepts the problem defined by the model as the problem he wishes to solve and proceeds accordingly. It is on this form of model usage that the natural sciences rest, and it is this form of model usage that is explicitly taught in technical and scientific education.

Newtonian mechanics, for example, exists as an abstract mathematical system of mechanics independent of any correspondence to reality. It also fits reality very well, at least in nonrelativistic situations, and is commonly used as a surrogate for real mechanics in trajectory calculations, for instance.

Another way in which a model can be used is as a *perspective* on the substantive problem, in the sense that two dimensional drawings or pictures serve to provide perspective on three dimensional objects. In using a model as a perspective, the analyst does not restrict himself to the structure of the model, but uses the model to organize and cue his deeper and fuller knowledge of the substantive problem, including aspects of the problem not represented in the model.

The use of models as perspectives is a familiar process. We do it all the time. We are doing it, in fact, whenever we interpret any

result derived from a model in a larger context than that represented by the model. We usually do it subconsciously, however, so that we don't think of it as a way of using models distinct from their use as surrogates.

The surrogate/perspective distinction is an important one in policy analysis. Most squishy problems are such that it is not possible to find an adequate quantitative surrogate for them. Any model used, therefore, should be viewed as a perspective, and the process of analysis, and particularly of interpretation, must take account of this fact.

The characteristic of being a surrogate or a perspective does not reside in the model alone, or even in the combination of model and problem. Rather, it resides in the head of the analyst--in the way he thinks about the problem. In method-oriented analyses, inadequate models are frequently used as pseudo-surrogates (with appropriate caveats) in the hope that the user can interpret the results in a broader context--i.e., can use them as perspectives. It seems likely that the quality of analysis could be improved by approaches to analysis that explicitly used the models as perspectives from the outset.

Understanding the use of models as perspectives requires an understanding of the nature of thought, or at least the distinction between verbal/logical/linear thought and nonverbal/intuitive/holistic thought. This is a distinction that man has recognized for centuries, although Western scientific thought has largely ignored it. Recent evidence, however, suggests not only that the distinction is a real one, but that the two thought processes take place in different sites in the brain.

These two thought processes are sometimes viewed as competitive--and perhaps as even incompatible. The verbal process is seen as the superior process, and the nonverbal as an inferior process that any right-thinking logical person would do well to ignore or suppress. This view is fundamentally in error, as the two processes are complementary rather than competitive. The nonverbal supports the verbal by suggesting ideas, paths of inquiry, etc., while the verbal serves to express, structure, and validate the nonverbal.

The verbal process plays two distinct roles in human thought. The first is as a thought process in its own right--as a manipulator and transmitter of ideas in the form of language. The second is as an interface between the nonverbal process and the outside world, as a way of expressing thoughts and understanding arrived at nonverbally.

These two roles are closely related to the surrogate/perspective distinction discussed earlier. At the purely verbal level, words serve as surrogates for the ideas they express, and those ideas can be understood in terms of the words alone. Used to express nonverbal understanding, however, words serve as a perspective on deeper more holistic ideas. The full meaning of those ideas often cannot be adequately understood by reference to the literal meaning of the words alone.

All this suggests the addition of another element to the conceptual model of analysis described above. That element is the analyst's internal *meta-model*, the understanding that he carries around inside his head of the problem and the tools, information, etc., he can bring to bear on it. This meta-model contains everything he knows about the problem, in some sense, but in a loose and ill-defined way. It is, in general, not capable of complete articulation in verbal (including mathematical) terms.

The process of analysis is a process of modifying and shaping portions of this meta-model, with the external constructs of formal problem and mathematical model serving as reflections of and perspectives on it. Formulation is the extraction of these constructs from the meta-model, and interpretation is reintegration of the results produced with the unarticulated knowledge within the meta-model.

In documenting his analysis and communicating it to others, the analyst should try to make his conclusions and their grounds as clear and explicit as possible. In mathematics, per se, and in rigorously quantifiable problems, this is done by making the grounds "free-standing" and independent of the meta-model. In the analysis of squishy problems, however, the nature of the problem precludes this and insures that the conclusions will remain judgmentally grounded, i.e., tied to the meta-model. To attempt to hide this fact in documentation by focusing on

the mathematical analysis and suppressing the role of formulation and interpretation is likely to distort the analysis and its conclusions.

Quantitative methodology has considerable potential as an aid to judgments and a source of insight into and understanding about squishy problems. That potential is diminished significantly when the methodology is looked to as a replacement for judgment and a source of objective knowledge, as it often is. The framework outlined here, primitive though it is, may allow a better understanding of this, by providing a broader perspective on the process of analysis and the role of considered subjective judgment in it.

CONTENTS

PREFACE	iii
SUMMARY	v
Section	
I. INTRODUCTION	1
II. POLICY ANALYSIS AND QUANTITATIVE METHODOLOGY	3
Policy Analysis	3
Quantitative Methodology and its Supporting Theory	5
Rigorously Quantifiable Problems	8
Other Applications	12
Method-Oriented Analysis	21
Quantificationism	26
Objectivity and Subjectivity	37
III. TOWARD A THEORY OF JUDGMENTAL ANALYSIS	41
Models	42
Surrogates and Perspectives	50
Right and Left Thinking	65
Verbal and Nonverbal Thought	75
Shamans and Systems Analysts	86
A Model of Judgmental Analysis	93
BIBLIOGRAPHY	105

I. INTRODUCTION

Recent decades have seen a tremendous expansion in the role and influence of quantitative methods and techniques in the governmental decisionmaking process. Operations research and related techniques found fertile fields of application during World War II, and their wartime successes were sufficient to guarantee them a continuing place in peacetime defense planning. Robert McNamara's tenure as Secretary of Defense, with its emphasis on systems analysis, program budgeting, cost-benefit analysis, etc., served further to institutionalize the use of quantitative methods, models, and techniques in defense planning and in the presentation of the results of that planning to Congress and the public. Concurrently, the use of quantitative methods and techniques spread to other departments of government, and to state and local governments as well. These methods increasingly pervade all areas of governmental planning and operations.

That use of these methods has, on occasion, proved fruitful and rewarding is without question. That it has, on occasion, proved misleading and led to unfortunate choices and consequences is also without question. The question of the net effect of quantitative methodology on the process of government is a complex one, probably not susceptible to definitive resolution. Some supporters of quantitative methodology and its application see it as the wave of the future--the natural extension of rationality and the scientific method to the process of government. Others, including this writer, are somewhat less sanguine.

This report provides a critical assessment of the usefulness of quantitative methodology as a policy analysis tool, and sketches the outlines of an approach to analysis and the application of quantitative methodology that might ameliorate, somewhat, the major weakness highlighted by the assessment.

That weakness (discussed in Sec. II) stems from the fact that the theory from which the methodology is drawn assumes the existence of a well-defined problem, and assumes that the model used in analysis is

an adequate and accurate representation of that problem. Based on that assumption, it prescribes methods for the logical analysis and solution of the problem defined by that model. Human judgment, per se, is not represented in or treated by the methodological theory.

Many of the problems encountered in policy analysis, however, are highly "squishy" and ill-defined. They may be equally well represented (though poorly, at best) by a number of distinct models. In such situations, the logical analysis of the model may be of relatively minor importance in determining the conclusions reached compared to the highly subjective activities of formulation--choosing a particular model from among the many possibilities available--and interpretation--deciding what the results mean in the light of the model chosen and the deficiencies present in it.

The absence of explicit treatment of these activities in the methodological theory often leads to their neglect in applications of the methodology. This neglect shows itself in a tendency to focus on calculation and the results of calculation and to ignore the issues of what is being calculated, and why. When this occurs, methodology becomes a poor and inadequate substitute for judgment, rather than a useful adjunct to it.

The approach to analysis that might ameliorate the problem (discussed in Sec. III) involves open and explicit acknowledgement of the role of subjective human judgment in any analysis of "squishy" and ill-defined problems. The use of mathematical models in analysis is a special case of the more general cognitive function of using models (of various types) to understand more complex phenomena. This report presents a discussion of that function and the beginning outlines of a theory of judgmental analysis.

II. POLICY ANALYSIS AND QUANTITATIVE METHODOLOGY

This assessment is not directed at quantitative methods and techniques *per se*. These are only tools, and one can no more fully understand the application of quantitative methodology in policy analysis by looking at the tools employed than one can fully understand the nature of gardening by examining rakes, hoes, and shovels. The subject of this assessment is the application of quantitative techniques and methods, with emphasis on the manner in which the tools are applied--the attitudes, beliefs, etc., of analysts applying the tools, and the approaches to problem solving and to the interpretations of results that those attitudes and beliefs engender. These issues of belief, approach, and interpretation are subtler, and far more difficult to deal with, than are the superficially straightforward and direct issues of method itself. They are, nonetheless, equally important in determining the value of quantitative methodology as a tool for policy analysis, and the worth of the knowledge and advice produced using it.

POLICY ANALYSIS

As used in this assessment, the term "policy analysis" should be broadly construed to connote the systematic examination or analysis of questions arising from, related to, faced by, etc., the governmental planning and decisionmaking process, conducted with the intention of affecting or contributing to that process. It thus includes forms of analysis such as force structure studies, program evaluations, etc., that might be excluded by a narrower use of the term. It includes studies done within the Air Staff, the Department of Defense, or other portions of the government, as well as by outside study agencies such as Rand when those studies are intended to affect governmental policy and operations.

The problems encountered in policy analysis run the gamut from well-defined to highly "squishy." A *well-defined* problem, in this sense, is one that can be given a clearcut, well-defined formulation, amenable to rigorous analysis. "How big a rocket is needed to put a

man on the moon?" for example, is such a problem. A "squishy" problem, on the other hand, is one with the property that any clear-cut, well-defined formulation of it will look like an unambiguous representation of the substantive problem only so long as we don't lean too hard on it, or question it too carefully or deeply. Thus, for example, the question "What strategic forces are needed to deter nuclear war with the Soviet Union?" is a very squishy question. The "Assured Destruction" formulation of that question, which measures the adequacy of a candidate strategic force in terms of its ability to inflict damage on the Soviet Union in a spasm exchange, is one way of turning it into an apparently well-defined problem. The appearance is only superficial, however, and evaporates rapidly when probed to any great extent. As a general rule, the well-defined problems encountered in policy analysis tend to be those concerned almost exclusively with technological or physical questions. Problems with any significant degree of behavioral or political content tend to be squishy. The more central the behavioral or political content becomes to the substantive issue, the squishier the problem is likely to be.

Policy analysis is carried on for a number of varied and sometimes conflicting objectives and purposes. Among these are the choice of a particular course of action to implement a predetermined policy, the advocacy or support of a particular position or point of view in the process of intra-governmental debate, and the search for illumination or insight into the nature of a particular issue or problem in order to better inform the policy- and decisionmakers who must deal with that problem. This last objective--of providing illumination and insight--is, in many ways, the most difficult and demanding. This is because in squishy problems there is no clearcut way of deciding how much illumination is enough, or measuring how much insight a particular analysis really provides. The present assessment is concerned primarily with this objective and with problems at the squishy end of the problem spectrum. For this reason, some of the criticisms made here may appear to be overstatements when related to better defined problems and less ambitious objectives.

The social sciences deal with many questions similar to those arising in policy analysis, but I will not be concerned with social

science methodology here. The distinction between the two, as far as I am concerned, is one of the objectives for which knowledge is sought. Social science reflects an academic interest in knowledge, while policy analysis reflects a more operational interest in supporting the ongoing functions of government. We may, therefore, impose problem-independent validity criteria on what we consider social science knowledge and simply recognize questions for which answers meeting those criteria cannot be found as being beyond the purview of scientific knowledge. Policy analysis, however, must deal with questions as they occur, and must provide understanding on the basis of existing information. Its criteria cannot be problem-independent since it does not have the choice of rejecting problems for which the criteria cannot be met.

QUANTITATIVE METHODOLOGY AND ITS SUPPORTING THEORY

As used in this report, the term "quantitative methodology" refers to the existing body of mathematical methods, computational techniques, etc., and supporting theory available to the policy analyst. This includes the theory of and tools derived from probability theory, game theory, statistical inference, econometrics, decision theory, computer simulation techniques, operations research, systems analysis, etc. These are overlapping terms with no clear dividing line between them. Each, in some sense, contains the others.

This methodology has an abstract existence and meaning--as a body of mathematical knowledge about ways of dealing with theoretical problems defined in and examined by the supporting theory--that is independent of, though related to, its application to substantive problems. The utility of the methodology as a tool for application to substantive problems depends on the existence of appropriate similarities between the theoretical and the substantive problems. In order to assess the application of the methodology to policy analysis, it will be useful to begin with a brief examination of the theory, per se, and of the nature of the problems it addresses and the solutions it provides to those problems.

The supporting theory can be classed as mathematics, per se, in the sense that it is a study of mathematical systems and models as abstractly

defined objects in their own right, independent of their relationships to real systems or real world substantive problems. The structure and behavior of mathematical models as such is determined fully by the defining premises and in no way by "real" considerations about any substantive process being modeled. Description of abstract models in "real world" terms, however, often serves as a useful heuristic device for understanding their behavior.*

Probability theory, for example, is the study of abstract mathematical structures called "probability spaces" or "measure spaces." Methodologies that derive from probability theory, such as statistical inference or decision theory, arise from the study of particular classes of problems that live on, i.e., are defined on or in terms of, those probability spaces. The problem of estimating the regression coefficients in a multivariate Gaussian process is a well-defined and meaningful abstract problem regardless of whether anything approximating a multivariate Gaussian process exists in the real world. Conversely, the fact that there exist "optimal" methods for estimating regression coefficients in such a process confers no necessary validity on those methods as tools for dealing with real problems involving data produced by something other than a multivariate Gaussian process.

Mathematical analysis, at the theoretical level, consists of the exploration of the structure and behavior of mathematical systems and models as logical consequences of the defining premises of those models. The results thus produced are grounded in the defining

* In this regard there is a significant difference worth noting between mathematicians engaged in mathematics, per se, and analysts engaged in the study of substantive problems in the way they relate mathematical models to the real world. The analyst concerned with the real world uses the model as an aid to understanding that world. He, thus, chooses his model to reflect relevant aspects of reality, and when the two diverge, will be likely to modify the model to correspond more closely with the appropriate reality. The mathematician studying the behavior of the abstract model, on the other hand, may be more interested in the behavior of the model than in the real world process he uses as a heuristic device to help him to identify with the model. When he talks of "tossing a coin," for example, as a heuristic description of a sequence of Bernoulli variables, his real interest is likely to be in the behavior of those abstract random variables--not the behavior of real coins.

premises and connected to them by a chain of logical inference. The product of such analysis is depicted schematically in Fig. 1. It consists of a mathematical model, the results derived from the model, and the chain of logic connecting the results to the model.

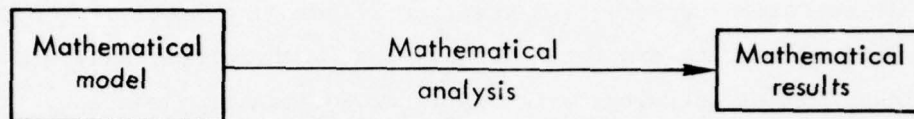


Fig. 1 — Mathematical analysis

The model is a well-defined mathematical object whose structure and behavior is logically determined by its defining premises, and, perhaps, by additional particularizing data (parameter values, boundary values, input data, etc.) that differentiate the particular instance or case under consideration from the broader class defined by the model. For example, a statistical model might describe the behavior of a statistical sampling process, with particularizing data describing the observed sample in a particular case. The results are precise statements about the further behavior or structure of the model, or about the theoretical problem defined by the model, that follow logically from its hypothesized structure. The grounds for the results consist of a chain of logical inference connecting the premises defining the model with the results derived from it. This chain may include (manual or computer) arithmetical computations since these are a form of logical inference.

The results produced by mathematical analysis are "objective" in the sense that they are grounded in a set of well-defined premises and logical inference from those premises. The structure shown in Fig. 1, if it is fully articulated, provides a representation of the results and the grounds on which they rest that might be said to be

"free-standing," in the sense that their validity can be determined on the basis of that structure alone, with no reference to or dependence on the competence or judgment of the analyst producing it. As such, they are subject to independent verification or repudiation by an outside critic on *logical* grounds alone. As we will see later, this a property not shared by many of the applications of quantitative methodology arising in policy analysis.

It sometimes appears from the free-standing, objective nature of mathematical results and their grounds that mathematical analysis is a purely logical activity, with little or no intuitive content. This is not the case. While the validity criteria of mathematics require the existence of a logical chain of proof that can be verified with no reference to intuition, initial discovery of that chain of proof is frequently a highly intuitive process. Good mathematicians develop a high degree of intuition about the mathematical models and theories with which they deal. They use this intuition in analyzing mathematical models and deriving results from them, particularly in producing original and creative results. The logical validity checks that mathematics imposes on those results, however, effectively "factor out" the role of intuition in the final product--producing the objective, free-standing product described earlier. This fact notwithstanding, creative mathematical analysis is a highly judgmental and intuitive process.

RIGOROUSLY QUANTIFIABLE PROBLEMS

The application of quantitative methodology to a substantive problem involves the use of a mathematical model as a simplified representation of the problem, or of some phenomenon important to the problem.* The substantive problem or phenomenon is analyzed through the analysis of the model. The nature of the conclusion reached, and the amount of credence and confidence that can be placed in it, depends in part on

*The terms "substantive problem" and "conclusion" or "substantive conclusion" will be used to refer to things of interest in the real world, and "model" and "results" to the corresponding entities at the mathematical level of analysis.

the results produced by mathematical analysis of the model. It also depends significantly on the relationship between the problem and the model--on what parts of the problem the model represents and how well, and on what parts of the problem the model distorts or fails to represent and how badly. Strictly speaking, this latter aspect of application--the effect of problem/model relationship--falls outside the scope of the theory on which the methodology is based since the theory deals with analysis of the problems defined by the models treated in the theory. This matters little when the problem/model fit is a good one, but can become important when the problem differs significantly from the model used to analyze it.

The simplest applications, from the viewpoint of the problem/model relationships involved, are those in which the structure and logic of the substantive problem are the same as (at least relative to the question being addressed) the structure and logic of the model being used. Such problems will be referred to here as "rigorously quantifiable." Rigorously quantifiable problems arise, for example, in the analysis of physical systems that obey well-understood and accepted scientific laws or in statistical experimentation in which the analyst may introduce randomness into his experiment to match the randomness in his statistical model. Analysis of rigorously quantifiable problems is methodologically pleasing because of the ease with which the theory transfers to application. The coincidence of structure between problem and model insures that the theoretically founded statement about the model will have valid analogs as statements about the problem.

Analysis of a rigorously quantifiable problem is depicted in Fig. 2. The analysis has two levels, a substantive level and a mathematical level. The analyst begins with the substantive problems on the substantive level. He adopts a mathematical model of that problem and analyzes the model to produce mathematical results. The model may not be spelled out in detail, particularly if it is commonly used and accepted for the type of problem he is dealing with (Euclidean geometry, for example, as a model in surveying or map making). It is understood to be there, nonetheless, and in principle, to be capable of rigorous specification. The results are then projected back up to the substantive level as substantive conclusions.

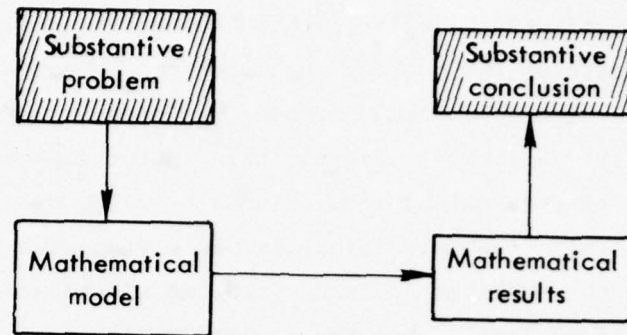


Fig. 2 — Analysis of a rigorously quantifiable problem

This is not to imply that the process of analysis is necessarily the straightforward, orderly process depicted here. Rather, when the process is completed, the end product can be sufficiently neatened up that it will appear as though it had been.

The calculation of the trajectory of an artillery shell is an example of a rigorously quantifiable problem. Newtonian mechanics is a mathematical model of "real mechanics" that has a well-defined structure independent of any relationship to reality. It seems to fit reality well (at least in nonrelativistic situations) and is commonly used as a surrogate for reality in analysis. To find the trajectory of an artillery shell using the Newtonian model of the problem, the analyst plugs in the particularizing data--weight of the projectile, force of the charge, elevation of the gun, wind, etc.--and solves the trajectory problem within the model. The analyst then accepts the result obtained as a substantive conclusion about the real trajectory of the real shell. This can be done, in fact, without ever consciously thinking about the fact that Newtonian mechanics is not "real mechanics" at all, but simply a mathematical abstraction that serves as a convenient surrogate.

A statistical experiment using random sampling is another example of a rigorously quantifiable problem. In this case, the analyst assures himself of a good model/problem agreement by structuring the problem to fit the model. He does this by designing a sampling

procedure (the problem) in a way that insures that it possesses the random character required by the statistical model that will be used to analyze the sample. The analyst can thus accept the model as a surrogate for the actual sampling procedure and accept the results derived from the model as conclusions about the substantive problem.

The validity of conclusions produced in this way depends on two factors: (1) the internal logical validity of the mathematical analysis and (2) the empirical validity of the linkages between model and substantive problem. The criteria for logical validity are the same as in the case of theoretical analysis discussed earlier. The difference between mathematical analysis and application lies in the fact that the premises defining the behavior of the model are no longer simply logical premises, but are assumptions about the behavior of the real world in the substantive problem. As such, they must be empirically validated. In problems dealing with physical systems obeying well-understood and accepted scientific laws, the validation is provided by the fact that the models used represent those scientific laws. In problems of statistical experimentation, empirical validity may be achieved by making the problem look like the model. In other cases, it may simply depend on the existence of long empirical experience with the process, sufficient to justify the treatment of the model as a surrogate for the process. This might be true, for example, of queuing models used to study the behavior of telephone exchanges.

If these logical and empirical validity criteria are met, then the conclusions reached are "objective" in the sense of being grounded in rigorous logic and objective fact. The structure of conclusions and grounds shown in Fig. 2 is, like that of Fig. 1, free-standing in the sense that its validity can be decided on the basis of the structure as it stands, without reference to or dependence on the competence or judgment of the analyst producing it. (This is true, at least, if the structure is fully articulated. It is true only in principle otherwise.) It is, thus, subject to verification or refutation by an independent critic on the grounds of the logic of the mathematical analysis within the model and the empirical connection between the model and the substantive problem.

Also, as in the case of the mathematical analysis, analysis of rigorously quantifiable problems conveys the false impression of being a strictly logical, rather than intuitive, activity. As in mathematics per se, intuition and good judgment on the part of the analyst may play a major role. But also as in mathematics, they are filtered out by the free-standing structure of verification and proof. Thus, while judgment may play a significant role in the formulation of the problem and in the conclusions reached, that judgment plays no part in the ultimate grounds for the validity of those conclusions.

OTHER APPLICATIONS

Rigorously quantifiable problems possess a structure sufficiently like that of the models used to represent them to justify the straightforward interpretation of results derived from the models as conclusions about the problems. If all problems to which quantitative methodology were applied were like this, there would be no need for concern about the limits of the methodology, or about applications not clearly justified by theory. That is not the case. Most of the problems to which quantitative methodology is applied in defense planning, and indeed in the analysis of public policy more generally, are not rigorously quantifiable problems. Rather, they span a spectrum from what might be called *reasonably quantifiable* problems at one end to highly *squishy* problems of dubious quantifiability at the other. Significant policy problems, in particular, frequently tend to lie much nearer the squishy end.

By a *reasonably quantifiable* problem, I mean one that possesses a fair degree of natural structure that seems to be reasonably represented by the mathematical model used to analyze the problem. Examples might include logistics problems such as inventory or maintenance scheduling problems. In many ways, analysis of reasonably quantifiable problems looks very much like the analysis of rigorously quantifiable problems. The major difference is that the requirements for empirical validity of the link between problem and model are less stringent in the case of reasonably quantifiable problems. The main criterion is that of reasonableness not of scientific validity.

A highly *squishy* problem, on the other hand, is unlikely to have any clear-cut, well-defined formulation that is both analytically tractable mathematically and that unambiguously captures the substantive problem. Clear verbal statements of the substantive problem may appear analytically intractable, while analytically tractable formulations of the problem are likely to significantly distort, or at least restrict, its substance. As a general rule, problems with significant behavioral or political content tend to be squishy. The more central that behavioral or political content becomes to the substantive issue, the squishier the problem is likely to be.

I want to look critically at this spectrum of problems and at the usefulness and the limitations of quantitative methodology across this spectrum. My emphasis will fall more heavily on the limitations of the methodology and on the squishy end of the problem spectrum--a combination often neglected in the advocacy literature for quantitative methodology. For this reason, some of what I have to say may appear overly critical when taken to apply to the application of quantitative methodology to reasonably quantifiable problems.

The extension of Fig. 2 to the application of quantitative methodology more generally is shown in Fig. 3. As in Fig. 2, there is a substantive level at the top, consisting of the substantive problem and the conclusion the analysis provides about that problem, and a mathematical level at the bottom, consisting of a mathematical model and the results obtained from the model using mathematical analysis. Unlike the rigorously quantifiable case, however, the model is rarely a direct match for the problem nor the conclusion a straightforward translation of the results. As the figure attempts to illustrate, both the substantive problem and conclusion may be far from sharply defined. The model itself will still be sharply defined, at least in principle, since that is the nature of mathematical models. This is true even if the model is not completely specified but is left partially implicit, since the mathematical results have rigorous meaning only with respect to a well-defined model. Suppose, for example, that a statistical test is applied and that the results are stated in terms of the "statistical significance" of the data. The use of a statistical model within which that term has meaning is implied, even if it is not explicitly stated.

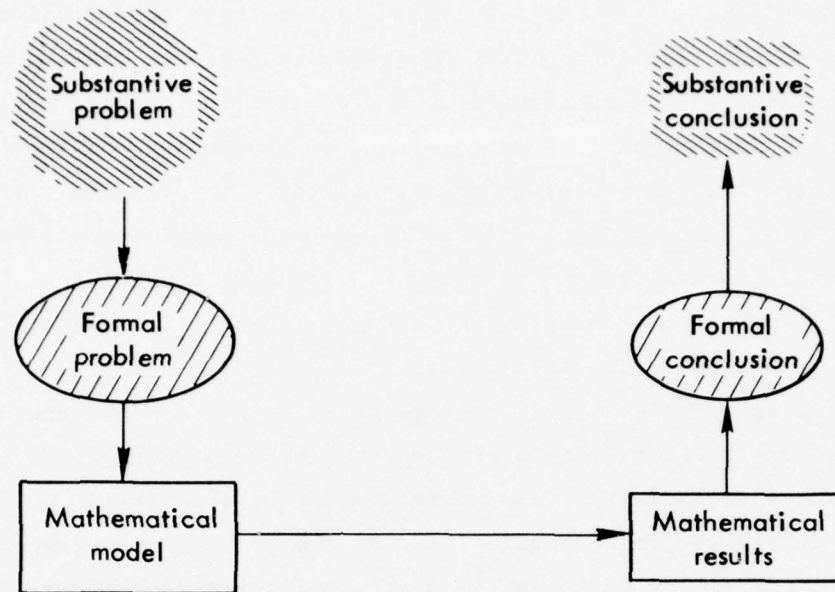


Fig. 3 — Application of quantitative methodology:
the general case

The relationship between the substantive problem and the model may be highly tenuous and ambiguous--particularly if the substantive problem itself is very squishy. The links between problem and model, and corresponding links between results and conclusion, may be far from clear. For this reason, it is useful to think of an intermediate *formal* level between the substantive and mathematical levels, serving to link the two. The formal problem links the substantive problem and the model by specifying what is being modeled in substantive terms. It delineates the parts of the problem addressed by the model, the assumptions being made, the parts left out, etc. It delineates the relationship of the model to the problem, perhaps describing or restating the model in substantive terms. The formal conclusion serves similarly as a link between the results and substantive conclusions--interpreting the results in a context more closely tied to the substantive problem and less to the model.

The formal problem and formal conclusion may not be explicitly identified and separated from either the substantive problem or model. They are, nonetheless, implicit in any application of quantitative methodology to a less than rigorously quantifiable problem and play an important role in determining the relevance and validity of the overall analysis. As Fig. 3 attempts to show, the structure of the formal level will typically be more sharply defined than that of the substantive level, but may still be considerably less precisely defined than the mathematical level.

If the substantive problem is reasonably quantifiable, as in the case, say, of an inventory stockage problem, then the formal and substantive problems may merge together and assume a structure almost identical with that of the model. As we move along the spectrum in the direction of increasing squishiness, however, the three become more clearly distinguishable, if not always more clearly distinguished.

Consider, for example, the comparison of alternative candidates for ground support aircraft to meet U.S. commitments to NATO, using a computer simulation of a NATO conflict. The substantive problem is one of determining which aircraft would better serve U.S. needs in the NATO theater. The model is the mathematical abstraction of a NATO/Warsaw Pact conflict embedded in the computer program used for the simulation. The formal problem in this case might be taken to be the evaluation of the aircraft in the specific conflict scenario of which that mathematical model is a model. The results produced by the computer simulation provide a measure of how well each of the abstractions of aircraft embodied in the model perform in the abstraction of war embodied in the model. These results can then be translated into a formal conclusion about how well the real aircraft would perform in the real conflict represented by the model. This translation of results into formal conclusion embodies judgments about the relationship between the model and the formal problem. Translating this formal conclusion into a substantive conclusion about the relative value of each of the aircraft involves additional judgments about the relevance of the particular conflict scenario considered in the formal problem to U.S. security interests in NATO.

Analysis can be thought of as consisting of the following three components:

1. *Formulation*: Formulation includes the choice of the formal problem and of the model to be used to represent that problem--the activity along the downward links on the left side of Fig. 3.
2. *Mathematical Analysis*: This occurs at the mathematical level, at the bottom of Fig. 3. It consists of analysis of and within the context defined by the model. This analysis produces results as mathematical statements about the model. It may also include purely technical questions of formulation and interpretation arising within the model itself, as opposed to those involving the relationship of the model with the formal and substantive problem.
3. *Interpretation*: This includes interpretation of the results as a formal conclusion, and of that conclusion as a substantive conclusion--the activity along the up links on the right side of Fig. 3.

In practice, these components are seldom separate, sharply defined activities occurring in sequence. Rather, they blur together as interwoven aspects of a complex and intellectual process. If they become separated, it is only when the results of that process are neatened up at the end. The formal problem and conclusion may not be explicitly identified as distinct from the model and the results or from the substantive problem and conclusion. Indeed, there is frequently no need for separate explicit identification, particularly in the case of reasonably quantifiable problems. Whether they are articulated as such or not, however, a formal problem and conclusion are tacitly assumed in the way results are presented and interpreted.

The validity of the conclusions reached depends on the logical validity of the mathematical analysis and on the validity of the linkages between problem and model. As is true with the cases considered earlier, the logical validity of the mathematical analysis can be

determined objectively, without reference to the subjective judgment of the analyst. This is not true, however, of the validity of the linkages between problem and model. In general, no objective standards exist by which these linkages can be validated. The thing that separates the rigorously quantifiable problems from those that are not, wherever they lie on the spectrum of squishiness, is the fact that the role of judgment along the links of formulation and interpretation can be factored out in the former but not in the latter problems. The conclusions reached at both the formal and substantive levels thus remain grounded in the subjective judgment of the analyst.

Formulation and interpretation are essentially subjective activities--requiring and depending on careful and considered judgment on the part of the analyst doing them. Formulation, from the substantive to the formal problem and from the formal problem to the model, is a process of taking away--of removing pieces to make the problem smaller and more analytically tractable. With complex and squishy problems particularly, this requires an intuitive understanding of the substantive problem as well as of the methodology being used. It may also involve some adding on, in the form of assumptions that are questionable on substantive grounds, but that make analysis easier. The assumption of statistical independence of various types of events, for example, falls in this category, as do assumptions about the "rationality" of political decisionmakers. Interpretation, conversely, involves putting things back--adding in the considerations removed to make the problem tractable and removing any distortions resulting from the added simplifying assumption.

The potential for distortion can be seen by considering Assured Destruction (AD) strategic force posture analysis--a form of analysis that has played a major role in U.S. strategic force planning over the last decade. The substantive problem is one of evaluating the deterrent capability of alternative strategic offensive forces. The mathematical model is a "worst case" strategic force exchange calculation in which Soviet strategic forces are applied to do maximum damage to U.S. strategic offensive forces, and surviving U.S. forces are then applied to Soviet value targets, i.e., cities. The results produced are estimates

of Soviet fatalities. If these fatalities exceed a specified AD level, the force posture under consideration is deemed to pass the AD test. If not, it is not.

One of the assumptions made in moving from substantive problem to model is the choice of the spasm war embodied in the model as the form of strategic conflict in which strategic forces will be tested. Does this represent a considered judgment that spasm war is the only strategic conflict worthy of substantive interest, or is it a simplifying assumption to make the problem more analytically tractable? This question is of more moment than it might initially appear since it has a major impact on the interpretation of the AD test as a force planning tool.

Should AD be viewed as one measure of adequacy or as the design goal for strategic forces? In other words, should AD be the predominant or only criterion by which we evaluate our strategic force posture, or should it be simply one measure of "How much is enough?" with additional evaluation and comparison of other force characteristics also relevant to any choice between alternatives? The AD criterion came into being in a period (the early 1960's) when the United States had an overwhelming strategic superiority. This superiority resided, moreover, in a force consisting of weapon systems developed and designed for different functions. It can be argued that the initial role of the AD criterion was more to say "How much is enough?" than to specify the design and composition of the strategic forces. Over time, however, it seems to have increasingly taken on the role of primary design goal for strategic force design, at least in the minds of many analysts. This is not the place to argue the substantive merits of this issue one way or the other, but simply to point out that it is, on the one hand, a critical substantive issue of U.S. defense policy, while on the other, an issue that can be decided by default in the formulation and interpretation of strategic force analyses.

A related but distinct question is that of what substantive assumptions about Soviet perception and behavior are reasonable. The simplest interpretation of the AD calculation in deterrence terms

assumes a Soviet leadership constantly motivated to go to war with the United States and deterred only as a result of explicit calculations of war outcomes. This leadership might be characterized as being advised by a "gnome in the basement" of the Kremlin who runs an AD calculation every Friday night and will advise attack on Saturday if expected Soviet fatalities fall below the AD level. The formal problem associated with this formulation is that of deterring that gnome. The AD calculation then has a direct interpretation as a formal conclusion--he will or he won't be deterred. This conclusion, however, requires careful additional interpretation in any broader context in which more realistic Soviet motivations and behavior are considered, e.g., in SALT considerations.

Alternatively, the formal problem can be viewed as one of ensuring that the outcome of total and unconstrained strategic warfare will appear unacceptable to any reasonable Soviet leader. With this formulation, the AD calculation is an ill-fitting but nonetheless useful model. The results produced by the calculation then require more careful interpretation as a formal conclusion, but at the same time, provide more insight into the larger substantive problem.

Formulation and interpretation are inherently judgmental in nature, and the conclusions reached are inextricably grounded in them. Hence, the conclusions themselves are judgmentally grounded. The degree of judgmental grounding, and its criticality to the validity of the conclusions, will vary with the problem. It will generally be greater for squishier problems. The *fact* of this judgmental grounding is inescapable, and should give lie to the claim that quantitative methods applied to squishy problems produce objective conclusions. The models and methods are tools, and their utility in complex problems depend on the skill and judgment of the analyst using them. To expect, as some do, that the methods themselves necessarily produce valid results is no more reasonable than to expect that a particular paint brush will produce a beautiful painting, or a particular knife a fine carving.

The theory from which the methodology is drawn is concerned with the application of the methodology to the problems treated by the

theory--problems specified as well-defined mathematical models with which the theory deals. Questions of how the methodology might be applied to problems outside that class (as most squishy problems are) are themselves beyond the scope of the theory and not considered by it. Some theories may address broader ranges of problems than others, but the inherent limitation remains. Many of the squishy problems arising in defense planning and analysis are far more complex than those dealt with in the theoretical constructs. The methodological literature tends to focus on example problems to which the methodology is applicable, with occasional examples of problems to which it clearly is not. The applicability (or usefulness) of the methodology in the gray areas in between (where most interesting applications fall) are treated largely with hand waving and reference to the exercise of good judgment by competent analysts and decisionmakers. Very little guidance is given, however, in the exercise of that judgment. The questions of formulation and interpretation are themselves complex, squishy, and ill-defined questions. It is, perhaps, easier and neater (from a methodological perspective) to pretend that they simply don't arise or that, when they do, they will somehow take care of themselves.

Another factor that contributes to the tendency to ignore the role of judgment in the application of quantitative methods is the fact that this role does not appear important in the area of inquiry in which the application of the methodology has proven most successful--the natural sciences. The problems addressed by the natural sciences are seen as rigorously quantifiable problems. The application of quantitative methodology in policy analysis is viewed by many of its proponents as an emulation of scientific principles and the scientific method. The fact that the methodology provides objective knowledge in science is seen as evidence that it should provide "objective" knowledge in other fields as well. This premise appears largely grounded in faith, however, and seems more a perversion of than an emulation of the principles of science.

This is not to imply that quantitative methodology is of no value in dealing with squishy problems. On the contrary, many examples exist in which the analysis of squishy problems using quantitative methodology

has provided significant insights or valuable solutions. This is usually due more to the skill and insight of the analyst, however, than to the power and rigor of the methodology per se. It stems not from the fact that he made calculations, but that he saw the *right* set of calculations to make, and was able to *interpret* them creatively. Nonetheless, the conventional dogma surrounding the application of quantitative methodology tends to attribute success to methodology per se, rather than to the wisdom and judgment of the analyst using it. It suggests that appropriate standards by which to judge quantitative methodology should focus on the technical quality of the mathematical analysis, rather than on the quality of the judgment associated with the formulation and interpretation of the analysis. These standards, in turn, encourage considerable amounts of sloppy analysis and confound the problem of separating the good from the bad.

METHOD-ORIENTED ANALYSIS

One result of inadequate attention to the role of judgment in the application of quantitative methodology is what might be called method-oriented analysis. In effect, the analyst pretends to be in the situation represented by Fig. 2, rather than that shown in Fig. 3. A model is chosen that is analytically tractable, and that can be tied to the substantive problem in a plausible way. The methodology is then applied in a straightforward manner, as if the model were a good fit to the problem and no judgmental issues were present in formulation or interpretation. The results are then given a straightforward interpretation in substantive form. The question of problem/model fit is finessed by the observation that conclusions follow from assumptions, and it is up to the readers, customers, etc., to decide whether or not they wish to accept the assumptions. The analyst is saying in effect, "Here's the problem I solved and my conclusion about it. Take it or leave it."

One major failing of the method-oriented approach is its almost exclusive emphasis on similarity between problems and models. It presumes, in effect, that if enough similarities can be found, then the methodology is applicable and will yield good results. This premise,

however, is true only if there do not exist enough differences to invalidate it. Unfortunately for the method-oriented analysts and their customers, however, the existence of many similarities does not preclude the existence of significant differences, as the following example illustrates.

Example: Regression Analysis

Consider a set of observations on two time series, X and Y , as shown in the scatter diagram in Fig. 4. Suppose that X is a policy variable, in the sense that it measures something subject to policy control, and that Y is a variable that measures something we would like to control. Suppose still further that common sense suggests that the relationship between the two is real, and not specious. What can be said on the basis of the data about the relationship between the variables and about the effect on Y of policy manipulation of X ?

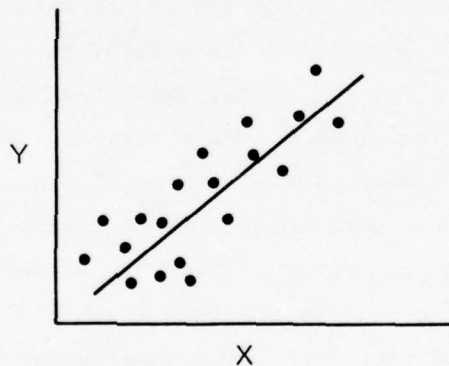


Fig. 4 — (X, Y) scatter diagram

For the method-oriented analyst whose method is regression analysis, the answer is clear. Fit a regression line to the data, as shown in Fig. 4, and use that line as a predictor of Y given X . Interpret the slope of the line as a measure of the incremental effect of changes in X on Y , and the correlation between observed and predicted values of Y (i.e., the corresponding points on the regression

line) as a measure of the expected accuracy of the prediction thus obtained. Using the method-oriented approach, then, it is possible to estimate the relationship between X and Y and to evaluate the affect on Y of policy manipulation of X on the basis of the available (X, Y) data--without worrying particularly about the judgmental questions of what X and Y really measure, what mechanisms really relate them, etc.

But how reasonable are the conclusions thus obtained? The answer to that question requires more knowledge about the problem than was specified above, or than a method-oriented analyst ususally takes into account in addressing it. Suppose, for example, that X is the mean annual depth of a reservoir, and that Y is the annual rainfall in the area. These variables have the characteristics outlined above. Because reservoir management policies can be changed, reservoir depth is a policy variable subject to policy manipulation. Annual rainfall is a variable we might be interested in controlling, and common sense clearly suggests that the relationship between rainfall and reservoir depth is nonspecious. If reservoir management policies have remained constant over time, then we might well see a strong correlation between rainfall and reservoir depth.

In spite of this, however, the conclusion suggested by the analysis--that we can decrease rainfall by draining water more rapidly from the reservoir and maintaining it at a lower level, and can increase rainfall by allowing more water to accumulate in the reservoir and maintaining it at a higher level--seems ludicrous. This is because the causal relationship assumed in the analysis--reservoir depth causes rainfall--runs counter to our common sense understanding of the relationship between the two--rainfall determines reservoir depth.

Because the nature of the mechanism connecting rainfall and reservoir depth is so self-evident, the speciousness of the conclusion is also self-evident. When there are more variables involved, the relationships between them less well understood, etc., the same kind of common sense test of validity may not exist. The plausibility of method-oriented analysis, coupled with a lack of any clear cut basis for disbelief, may then appear as strong grounds for acceptance of the proffered conclusions. A number of analyses of the Viet Nam conflict

were conducted in the late 1960's that rested largely on arguments closely parallel to these. This example, however, shows how specious such conclusions can be in the absence of careful evaluation of the fit between problem and model and interpretation of results in light of that fit, i.e., of questions of substantive formulation and interpretation shown in Fig. 3.

An equation of the type found in regression analysis, say, such as,

$$Y = aX + b + e, \quad (1)$$

where a and b are constants and e is an error or disturbance term, can be assigned two distinct types of meanings. The first is a *statistical* meaning--the equation describes a statistical relationship between X and Y , viewed as the joint output of some otherwise unspecified process. The statistical meaning refers only to the joint distribution of the variables, and not to the *causal* mechanism relating them. The equation can also be given a *causal* meaning in the sense that it describes a system in which the output Y is produced by inputs X and e . There is no necessary connection between the two meanings. To see this, note that the statistical meaning of the equation is unchanged if the equation is transformed into

$$X = a'Y + b' + e', \quad (2)$$

where a' , b' , e' are appropriate linear combinations of a , b , and e . The causal meaning, however, is reversed.

Relative to the mathematical systems that such equations are conventionally understood to describe--multivariate Gaussian processes--there is no need for distinction between the statistical and causal meanings. Statistically, the equation describes the joint distribution of random variables X and Y . Causally, an (X, Y) pair having this distribution can be thought of as being obtained by sampling X from its marginal distribution, then sampling Y from its conditional distribution given the observed value of X . This mathematical system is equally adaptable to the alternative causal interpretation provided by Eq. (2)--simply reverse the order of sampling.

When equations such as Eq. (1) are used as models of natural systems, however, a distinction between these two meanings is necessary. If X is reservoir depth and Y is rainfall, then Eq. (1) is a meaningful description of a statistical relationship between them, but not of a causal relationship. Statistical techniques such as tests of significance apply to the statistical meanings of equations such as Eq. (1). Hypotheses arising in substantive policy problems, however, are more often causal than statistical in nature. An equation such as Eq. (1) (or its more complex multivariate extensions) can be used to describe a natural system, and a statistical test can be applied to a causal hypothesis about that system. The validity of that test depends on the validity of the identification of the causal and statistical meanings of the equation relative to that natural system. That identification requires separate distinct judgment made outside the context of the test itself. If that judgment is not made, the risk of conclusions equivalent to "reservoir depth causes rainfall" is high.*

This type of error can result from confusion in the meaning of words, as well as equations. To say that data are "statistically significant" says only that the data do not appear to have been produced by the mathematical system specified as the model for the "null hypothesis." Leave aside for the moment questions of how closely that system parallels the state of the natural system described by the null hypothesis--e.g., questions such as the randomness of sampling or the linearity of relationships. Even apart from such questions, there is no *a priori* reason to equate "statistically significant" with "operationally significant" without careful ad hoc judgment relative to the question being addressed. Yet this distinction is frequently blurred.

Another oft confused term is "independent variable." The following are among the meanings it seems to have:

1. On the right of the equal sign.
2. A cause in a cause-effect pair.

* For additional discussion of this point, see Ralph E. Strauch, *The Fallacy of Causal Analysis*, The Rand Corporation, P-4618, April 1971.

3. Chosen first in the sequential selection of a multivariate sample point (the process described above).
4. A variable over which we have some policy control, e.g., we can control reservoir depth by changing our reservoir management policies.

These meanings are all distinct. They may or may not coincide in a particular situation. Assuming a greater degree of coincidence than is reasonable between them is a major cause of fallacious statistical inference. Particularly pernicious is the oft made (usually implicitly) assumption that 1 and 4 are sufficient to guarantee 2. It is this assumption, for example, that leads to the inference that reservoir depth causes rainfall and to other less obvious but equally specious inferences in more complex situations.

Quantitative analysis is largely a process of manipulating symbols--numbers, equations, mathematical statements, etc. If the equations, statements, etc., have different meanings in different parts of the process, then the validity of the manipulations made will be critically dependent on the coincidence of those meanings in the particular problem under analysis. If they fail to coincide, the conclusions may rest on the same illogic and confusion of meanings that, when applied to the word "cardinal," would support the conclusion that the St. Louis ball team belongs to the hierarchy of the Roman Catholic church. The symbols involved may include large amounts of multivariate numerical data, rather than a few simple natural language sentences, and the manipulations consist of detailed computer analysis using complex multiple regression programs rather than simple verbal syllogism. The same underlying fallacy can occur at either level. In neither case, moreover, is the issue of possible confusion of meanings one that can be settled within the context of the formal analysis--the symbols being manipulated and the manipulations being made. It can be settled only through separate distinct judgment made apart from that manipulation.

QUANTIFICATIONISM

In spite of its limitations, there is a contemporary school of thought that advocates the method-oriented application of quantitative

methodology to a broad range of behavioral and political problems arising in policy analysis. It seems to hold that the issues of formulation and interpretation are either nonexistent or will somehow be automatically taken care of by the competence and good judgment of analysts and their customers, even when never *explicitly* surfaced and considered. This school of thought, which will be referred to here as *quantificationism*, seems to hold that quantification is a positive value per se, i.e., that for most questions, a quantitative answer is *a priori* better than a qualitative one. Its more extreme proponents seem to hold the view that any problem of substance is reducible to a form that admits a quantitative solution and that no problem or issue can be adequately understood until it is reduced to such a form. Its more moderate adherents might reject this extreme position. They would nonetheless argue that quantification almost always improves understanding and decisionmaking, and that as a result, quantitative analyses should be conducted whenever possible.*

Quantificationism comes in a variety of forms, but its salient features are an emphasis on methodology per se, i.e., on models and techniques and the results that they produce, and the belief that this emphasis is appropriate in the application of quantitative methodology to squishy problems--that it captures the scientific method and provides a scientific approach to such problems. It appears to do this because it appears to emulate the reductionism inherent in the physical sciences.[†] It can be argued, however, that this reductionism should

*The reader who is so inclined should have no trouble quibbling with the details of my characterization of quantificationism. I could easily do so myself. The relevant question, however, is whether a doctrine like quantificationism is operative in our society today--not whether or not I have managed to define it precisely. A colleague once suggested that quantificationism must be a strawman, because no competent analyst would defend quantificationism as I have defined it, or would admit to the practice. That argument assumes a consistency of belief and action that few men achieve. It is hard to find a defender of hypocrisy or self-deception, either. Nonetheless, we all engage in them on occasion.

[†]*Reductionism*, in this context, refers to the thinking style that treats a phenomenon of interest as an isolated system, develops a quantitative model for that isolated system, and uses that model as a surrogate for the phenomenon in analyzing or answering questions about the phenomenon.

not be considered as the central core of the approach to knowing inherent in the physical sciences, but rather as a characteristic of that approach as applied to the particular problems that the physical sciences address. In attempting to learn from science about how to understand other classes of problems, therefore, we must question whether or not that particular characteristic of science is one that should transfer to this new class of problems.

Rationales for quantificationism are of two basic types--literal and pragmatic. The literal rationale derives from the assumption that the methodology does in reality what it claims to do in theory. The pragmatic rationales, on the other hand, acknowledge the limitations of the literal rationale, but argue that the method-oriented application of quantitative methodologies is nonetheless a valid approach to analysis.

Decision theory, for example, deals with choice in the face of uncertainty. In particular, it deals with the problem facing a decisionmaker who must choose between alternative courses of action in the face of uncertainty about what outcome will result from the action he chooses. It prescribes the way in which the decisionmaker should integrate his uncertainties and his preferences among (or utilities for) outcomes, in order to select the best decision in the light of those uncertainties and preferences. The elements of the decision problem dealt with by decision theory--the decisionmaker's alternatives, uncertainties, outcomes, preferences--are elements that can be found in most problems of policy choice. Hence, decision theory appears to provide a useful methodology for addressing and dealing with such problems. Indeed, many quantitative analyses are of this form--the model used is a quantitative decision problem, and the results of the analysis prescribe the optimal choice within that problem. The literal rationale for such analyses derives from equating the words used in the theory--decisionmaker, uncertainty, preference, etc.--with the same words in the substantive problem. Thus, "uncertainty" in the substantive problem is assumed to be adequately captured by "uncertainty" in the model, the set of alternatives in each are assumed to be the same, etc. It follows logically that the choice prescribed by analysis of the model is optimal in the substantive problem.

In cases where this literal identification of problem and model can be justified, the literal rationale provides strong grounds of support for the analysis and conclusions. These cases, however, are more the exception than the rule, particularly among problems requiring any significant degree of creative insight. In most problems, the validity of literal problem/model identification will be somewhere between questionable and obviously specious. In these cases, the literal rationale is invalid, and if the methodology is to be justified, a different rationale must be found. The pragmatic rationale says, in effect, that even if the methodology doesn't do what it claims, it's still a useful aid to decisionmaking, or to understanding squishy problems. The reasons why this is claimed to be true vary, ranging from "You wouldn't want to trust judgment alone, would you?" to arguments that quantitative analysis helps lay bare the structure of the problem and provides decisionmakers with a better understanding than they would have otherwise had. There are, in effect, different pragmatic rationales applicable in different degrees to different problems. None are supportable or refutable on strictly logical grounds.

The Genie in the Eight Ball

The structure of these rationales can perhaps be seen more clearly if we set quantitative methodology aside for a minute and examine another decision aid or source of knowledge--the Genie in the Eight Ball. Most five-and-dime and magic stores sell a large plastic eight ball with a small window on the bottom, filled with a black liquid. Inside is a polyhedral float with inscriptions such as "yes," "no," "ask me again," "better not," etc., on the faces. To get advice from the eight ball, you phrase a question that can be answered yes or no, turn the ball over, and see what answer floats up to the window. One rationale that might be given for the efficacy of this procedure is that the answer is determined by a Genie who lives in the eight ball, and who talks directly with the gods who control the world.

Now, suppose I have a friend who advocates using the Genie in the Eight Ball to deal with squishy problems. Let's examine the rationales he might put forth to support that procedure, and the parallels that exist in rationales for quantificationism.

My friend's basic rationale might well be that the Genie speaks directly to the gods, and the gods determine how the world works. Asking the Genie is thus justified by the fact that he provides a direct line to the gods. This theological argument corresponds to the literal rationale for quantificationism, in the sense that it is obtained by interpreting the description of the methodology literally. Occasional failures of the Genie to provide good advice do not negate this argument. They can be explained by the fact that the Genie is capricious, that he may not always understand the gods, or that the gods may change their minds after they've talked to him. If my friend truly believes, there is likely to be little I can say to convince him otherwise, and if I am truly skeptical, there will be little he can do to convert me.

The major difference between this theological rationale for the Genie and the literal argument for quantificationism lies in the fact that I can accept or reject the former on a one time problem-independent basis, while the latter requires an *ad hoc* determination based on the nature of the question. In other words, I accept or reject the existence of the Genie independently of the question I asked him to answer for me. I cannot make the same problem-independent decision about the validity of mathematical models of real phenomena, however, since I can find some instances in which models are clearly valid and others in which they are clearly invalid. The question of validity for a particular problem/model pair, therefore, is one that must be dealt with on an *ad hoc* basis. While this complicates the problem of accepting or rejecting quantitative analysis, it does suggest rejection of the literal rationale for quantificationism. Too many counter-examples exist for universal acceptance of that rationale.

But what of the pragmatic? Consider my friend with the eight ball. If I point out to him that I don't believe in the Genie and, hence, think the eight ball is a lousy decisionmaking device, he might respond that he doesn't believe in the Genie either. He might say, in fact, that anyone would be silly to really believe in the Genie, but that regardless of whether the Genie really exists or not, the eight ball is a useful pragmatic decisionmaking tool. Among the reasons that he might cite are the following.

Managerial. He helps make arbitrary decisions. Many of the decisions most people or organizations face are largely arbitrary in the sense that the critical choices have been made in selecting the alternatives. The final choice doesn't really matter much, at least so far as can be told on the basis of current knowledge. The choice of where I go to lunch, for example, is in this category. As long as I am reasonable in the set of alternatives I consider--i.e., I pick places nearby, within my price range, and with reasonably sanitary conditions--my actual choice is irrelevant, and it would be a waste of time to agonize over the decision. If I have trouble making arbitrary decisions and need some sort of rationale for doing so, asking the Genie and accepting his advice can be a very useful procedure.

Empirical. His advice has been good in the past. My friend might argue that the Genie has shown himself to be useful and reliable in the past, as demonstrated by the worth of the advice he has provided. If the decision of where to go for lunch is not recognized as essentially arbitrary, for example, the fact that never once had following the Genie's advice led to food poisoning (or even indigestion) might be cited as evidence of his efficacy. My friend might also be able to cite a few more significant decisions in which the Genie had played a major role. Suppose, for example, that he had recently considered the purchase of either a station wagon or a motorcycle, and on the Genie's advice, had purchased the station wagon. If he has a wife, four kids, and a large dog, and he likes to take weekend family outings, he might point out that the motorcycle would have clearly been an inferior choice. The Genie's advocacy of the station wagon thus appears to provide further evidence of his efficacy and reasonableness.

Catalytic. He jogs the thought process. My friend could also argue that he knows better than to take the Genie's advice without question, but that discussing things with the Genie is a useful catalyst in clarifying issues and in identifying important factors that might otherwise go unnoticed. In effect, he might say that it's not the answer that matters, but the way he asks the questions. His seeking advice on the motorcycle-station wagon decision, for example, may have involved a long and drawn out dialogue on the relative merits of

each, rather than a simple yes-no question of "Should I buy a motorcycle?"

Ritualistic. In further support of any of these rationales my friend might argue that the Genie is a useful social convention, a ritual that people use even though they know better, in the same way we anthropomorphize "Mother Nature" or "Jack Frost." He could argue, in effect, "Of course I don't believe in the Genie, but he provides a useful terminology and a useful framework for describing how I do things."

Contingent. The answers the eight ball provide are contingent on the existence of the Genie, and it's up to the user to decide whether he is real or not. Suppose that, in addition to (or perhaps instead of) using the Genie to make his own decisions, my friend makes a living as a fortune teller, providing advice from the eight ball to others. He then need not evaluate his own belief in the Genie. Instead, he might argue "All I do is tell the people what the Genie says, and it's up to them to decide whether they believe in him or not and to accept or reject his advice accordingly. I lay no claim to knowledge, but rather to the ability to turn the eight ball over and read the answer that comes up. What I do is clear, and it's up to the customer to decide whether or not it is useful to him."

Unlike the theological argument, these are not rationales that can be universally accepted or rejected on a problem-independent basis. They all have some merit, depending on the question asked, the alternative forms of inquiry available, and other contextual factors. If my friend would find himself in a constant quandry and turmoil over where to eat lunch or what kind of gasoline to buy, for example, the Genie may serve a very useful function in providing him a means of dealing with that type of question. If he feels that the eight ball really has worked in the past, then, in some sense, that justifies using it in the future. This is particularly true if the Genie is in fact a good catalyst to his thinking--if conversations with the Genie truly served to peel back the layers of my friend's consciousness and expose information, ideas, etc., that he might otherwise fail to consider. If he makes his living with the eight ball, then the contingent

rationale is particularly cogent. In using that rationale, however, it is important to distinguish between belief in the efficacy of the method and economic self-interest. If my friend wants to truly understand the validity of what he does, he must be careful to avoid overselling himself on the existence of the Genie in order to rationalize his livelihood. If he wants to consider himself an ethical fortune-teller, then it is incumbent on him to exercise a great deal of care in insuring that he doesn't oversell his method to his customers, but that he is as careful in defining its limits as he is in touting its advantages.

These rationales all have parallels among the rationales for quantificationism. Moreover, because quantitative analysis is a far more complex methodology than the Genie in the Eight Ball, these parallel rationales are correspondingly more complex.

Many of the decisions faced by government, especially at fine grain levels of detail, are of the largely arbitrary type. Final budget allocations, once reasonable limits have been decided on, for example, are in this category, as are most choices between largely similar but competing programs. This is particularly true when a fair margin of uncertainty exists about the precise utility of the program or about the incremental return from additional investment. In such cases, well-defined repeatable procedures to make precise choices among reasonable alternatives are needed. Quantitative procedures satisfy this need very well.

Many of the empirical arguments for quantificationism rest heavily on this managerial justification--"Look how much smoother things have run since we introduced program budgeting," etc. Quantificationist techniques and procedures are often highly efficient as management tools, largely because of the extent to which they simplify problems and provide repeatable mechanical decision procedures. The very characteristics, however, that make any methodology (quantitative methodology or the Genie) a useful managerial device for routine decision-making make it a risky device for catalytic mind jogging. The risk occurs because the terminology used, etc., make it quite easy to attribute validity to the mechanical application of the methodology

rather than to its use as a catalyst. Once this is done, it is easy to slip into the habit of applying it mechanically and accepting the answers thus obtained as justified because of the way they were obtained. In spite of the fact that my friend found the Genie a useful device for exploring the relative advantages of the motorcycle and the station wagon, for example, he is deluding himself if he attributes the wisdom of his decision to buy the stationwagon to the Genie. In so doing, moreover, he increases the likelihood that he will misuse the Genie next time, by listening to the Genie and ignoring his own good sense when the two conflict.

The ritualistic and contingent rationales are particularly pernicious in this regard, because they encourage sloppiness about just what is being analyzed and just what the analysis shows, and a sense of complacency about the potential impact of that sloppiness. They also serve as "cop outs," relieving the analyst of responsibility for the substantive interpretation of his conclusions, and transferring that responsibility to "them," "everybody," or the customer for or reader of his analysis. The ritualistic rationale takes such forms as "This formulation leaves a lot out, but it is the one everybody uses so it will do for now," or "Everybody uses that terminology and understands its limitations," etc. Contingent rationales occur in justifications such as "It's not up to me to decide if the assumptions are reasonable, but only to insure that the conclusions follow," etc. These rationales seem to assume that the good judgment of all involved is sufficient to overcome the weaknesses and inadequacies of literal interpretations of method-oriented analysis. This assumption is in direct contradiction to that portion of the literal justification which argues that explicit systematic analysis is necessary because of the inadequacies of subjective judgment.

An Image of Analysis

There is a prominent contemporary image of the quantitative analyst as a purveyor of objective fact based on logic and hard data. He takes no personal responsibility for his conclusions, since they are not of his making, but are inherent in the nature of things. All he

has done is uncovered them and made them visible for all to see. This is conveyed by the fact that he never talks about what *he* thinks, but restricts himself to what "the analysis shows" or "is clear from the data." In the sense that he is perceived as not personally involved with his conclusions, he is like the natural scientist, or perhaps the priest who serves only as a conduit to the gods. He distains the "merely qualitative" and often speaks perjoratively of "subjective judgment."

Now very few people (there are some) "really believe" in this image, in the sense of accepting it without qualification after careful reflection. This is the image, nonetheless, that serves to define the operational role of the quantitative policy analyst in contemporary society. He pretends that it is true, and his customers pretend that it is true, even though all know better.

This results in a set of conventions that define how the analyst will perform and report on his analysis, and how the customer will react to it. The analyst chooses a model he knows how to analyze, analyzes it as though it were an adequate representation of his problem, and interprets his results accordingly. He reports these conclusions and describes enough of the model and the mathematical analysis to show that he performed the ritual correctly. This description usually contains a listing of "assumptions" from which the conclusions are said to follow. It frequently contains the disclaimer that the evaluation of the assumptions is the responsibility of the customer, not the analyst. It may also contain the disclaimer that the analysis addresses only a part of the problem with which the customer has to deal, and additional factors must be taken into account by the customer in acting on the analysis. In spite of these disclaimers, however, the conclusions are frequently stated in positive and unequivocal terms.

These conventions allow for the production of a great deal of good analysis, since they give competent and insightful analysts much leeway within which to operate. They also allow the production of much bad analysis, since they focus attention primarily on the technical correctness of the mathematical analysis rather than on its

substantive validity or relevance. This, in turn, encourages sloppy method-oriented analysis of little substantive value.

The government policy- or decisionmaker who is the customer for the analysis also has a role to play in support of this image. He pretends he believes in the image, while at the same time, independently judging the worth of the analysis presented to him. He must reject the bad and use (and perhaps even learn from) the good. The hedges provided by the analyst, in the form of assumptions the customer can disagree with or "other factors" not included in the analysis that he must consider, allow him to do this without seriously questioning or coming into direct conflict with the image of the analyst as a giver of fact.

Everyone goes along with the charade because everyone else does, and it seems to be a relatively harmless charade as long as everyone understands it.

But is it? That itself is a very squishy question. It involves the assessment of how well people really perform when doing one thing while pretending to do another, neither of which is well defined in any formal objective sense. There is no objective way, really, to decide. A number of arguments, however, suggest that we would get more out of our policy analysis if we were more honest about the difficulties of the task, and about the importance of careful subjective judgment in coming to grips with and coping with those difficulties.

There is no way of really being sure how well people really understand the nature of the charade. The amount of bad analysis done by seemingly competent and sincere analysts would suggest that they don't all fully understand it. At the same time, some of the arguments made by policymakers that seem to be based on or swayed by bad analysis would suggest that they don't all really understand it fully either.

Even if they do, role-playing gets in the way. It tends to encourage bad or sloppy analysis by analysts who should, and do, know better, because that's what they see the system as expecting of them. It tends to encourage sloppy evaluation and use by the customers of analysis, particularly when the analysis appears to support a position toward which they feel well disposed.

The charade discourages thoughtful systematic consideration of nonquantitative factors. It does this by defining the ground rules for debate--the terms in which issues are to be framed and decided. It directs attention and effort to the quantitative aspects of basically unquantifiable problems and tends to force debate into that mold. It becomes, somehow, too inconvenient, too cumbersome, etc., to systematically address the aspects of the question that don't fit in the quantitative mold.

Over time, this introduces institutional distortions, in the sense that people begin to accept simplistic characterizations of problems as the institutional way of looking at them, even though they, individually, know better. What happens eventually, then, is that the institution as a whole settles for and operates with characterizations of its problems that most people in it know is inadequate.*

All of these factors combine to suppress the positive potential that systematic analysis has to provide in understanding and coping with complex and squishy policy problems. On balance, then, the contention that the charade is harmless is questionable. It has a positive potential for harm and this potential seems sufficiently great to make efforts to find alternatives worthwhile.

OBJECTIVITY AND SUBJECTIVITY

One of the underlying imperatives toward quantificationism is a desire for "objectivity," without an adequate understanding of what "objectivity" means in the context of analysis of very squishy problems and of the tradeoffs involved in achieving different kinds of objectivity. The natural sciences are seen as objective--as unbiased, quantitative, logical, and free-standing in the sense discussed earlier. These characteristics are equated with objectivity, and the positive valuation attached to them as attributes of science is transferred to

*The effects of this phenomenon in strategic defense planning are discussed in my "Winners and Losers: A Conceptual Barrier in Our Strategic Thinking," *Air University Review*, Vol. XXIII, July-August 1972, pp.33-44. Also available as P-4769, The Rand Corporation, June 1972.

the characteristics themselves. Simultaneously, the converse characteristics--biased, qualitative, intuitive, and judgmental--are denigrated and labeled "subjective." "Objectivity" is seen as good, desirable, something to be emulated while "subjectivity" is viewed as negative, perjorative, something to be avoided. This labeling identifies qualities (such as "quantitative" and "unbiased") as always occurring together, and hence, equivalent. They do seem to occur together in the natural sciences, but this will not necessarily be true in the kinds of squishy problems encountered in policy analysis.

It is important, therefore, to distinguish between different meanings of "objective" and to identify situations in which those meanings may coincide or be contradictory. When different forms of objectivity are not simultaneously achievable, we must choose between them. In some circumstances, the most "objective" analysis may be analysis grounded in a high degree of carefully considered subjective judgment.

The term "*objective*" has two distinct meanings relative to analysis and the conclusions produced by analysis. The first is what I've referred to earlier as "*free-standing*"--detached from the analyst and residing in a literal statement of the conclusion and its grounds. Objectivity, in the free-standing sense, is a characteristic of a conclusion, or product of analysis, and occurs when the conclusion is grounded entirely in logic and intersubjectively verifiable fact. A second and distinct meaning of "objective" is "*unprejudiced*"--coming to the problem without preconceived ideas about what the conclusion should be, trying to see things as they really are rather than as the analyst would like them to be, etc. Objectivity in the unprejudiced sense is more a characteristic of the analyst and of the process of analysis than of the conclusion per se.

Corresponding to these two forms of objectivity are complementary forms of subjectivity--judgmental and prejudiced. *Subjectivity*, in the *judgmental* sense, is a property of the conclusion and the grounds--the property that the conclusion remains grounded, at least in part, in the subjective knowledge and judgment of the analyst. The grounds reside with the analyst in ways not fully articulated in the language describing them, or fully representable in terms of logic and objective

fact. Subjectivity, in the *prejudiced* sense, is a property of the analyst or the process of analysis he brings to bear. It means a coming to the problem with a preconceived bias about the nature of the problem or about the conclusion to be reached.

The logical compatibilities between these definitions are shown in Fig. 5. "Free-standing" and "judgmental" are mutually exclusive properties. Hence, they are logically incompatible. Similarly,

Objectivity	Subjectivity	
	Judgmental	Prejudiced
Free-standing	Incompatible	Compatible
Unprejudiced	Compatible	Incompatible

Fig. 5 — Logical compatibilities between forms of objectivity and subjectivity

"prejudiced" and "unprejudiced" are mutually exclusive and incompatible. There is, however, no *logical* incompatibility between a prejudiced analyst and free-standing knowledge or between an unprejudiced analyst and a judgmental conclusion.

These two forms of objectivity seem to coincide in the natural sciences--unprejudiced inquiry leads to free-standing knowledge. This coincidence is a property of a class of problems that the natural sciences address, however. It is not one that we have any strong a priori reason to expect in other areas of inquiry or analysis. To insist on the production of conclusions that superficially *appear* free-standing without first determining that the conclusion sought can be reasonably expected to have that property is itself a severe form of prejudice. What it is likely to produce (and often does in method-oriented quantitative analysis of squishy problems) is superficially free-standing pseudo-knowledge of questionable relationship to the substantive problem. This pseudo-knowledge may be useful, but if it

is, it is only through careful judgmental interpretation, in combination with additional judgmental knowledge of the problem possessed by the interpreter.

Much as we might wish it so, there is no a priori guarantee that the two meanings of objectivity will necessarily coincide in policy analysis. Indeed, there is good reason to expect that they will not, since many of the problems encountered possess a political/social/behavioral content of the type man seems best able to understand and cope with judgmentally. In such problems, we cannot "see things as they really are" while restricting our attention to the quantifiable or analytically tractable aspect of the problem. We cannot, therefore, have both forms of objectivity, and we must choose between them. I believe the choice should be in favor of unprejudiced objectivity--of attempting to understand the problems as thoroughly as possible--acknowledging that this means acceptance of judgmentally grounded conclusions in many cases. To insist on the superficial appearance of free-standing objectivity--as the quantificationist does--is itself a severe form of subjective prejudice that can significantly distort the analysis and the conclusions reached.

Again, this is not to imply that quantitative methods have no role in the analysis of squishy problems. They do. That role should be as an adjunct to and in support of carefully considered subjective judgment, however, rather than as an attempted replacement for judgment. For reasons outlined earlier, current theories from which quantitative methodology derive provide inadequate guidance as to how this might be done. The remainder of this paper will be devoted to sketching some preliminary thoughts concerning a theory of judgmental application of quantitative methods.

III. TOWARD A THEORY OF JUDGMENTAL ANALYSIS

One of the major limitations on and sources of abuse of quantitative methodology as a tool for the analysis of squishy problems is the fact that the theories that support the methodology and guide its application do not admit a role for subjective human judgment. This is because they deal with fully structured problems (the problems defined by the models employed) that are sufficiently specified and well defined to insure that the results sought follow logically from the structure of the problem. Unfortunately, few of the problems arising in defense planning or other areas of policy analysis are like that. Most are, instead, squishy to varying degrees, with an ambiguous natural structure that can be plausibly modeled in a number of ways, none of which is precise or empirically verifiable. The components of analysis that figure most prominently in such problems, formulation and interpretation, are outside the scope of existing quantitative theory.

But these are precisely the components in which judgment plays an overriding role. Formulation involves judgments about what the problem is, how it should be structured to make it analyzable, and what model can be usefully employed in analysis, in spite of the inadequacies of the model and the distortions it introduces. Interpretation involves judgments about what the results obtained really say about the substantive problem in the light of those inadequacies. The overall conclusions depend heavily on, and in fact, are ultimately grounded in, those judgments.

When good analysts do good analysis, it is because they make those judgments well. They do this, however, distinct from, and in some cases, in spite of, the existing methodological theory and conventional dogma surrounding the application of that theory, rather than because of it. This suggests a need for an expanded theoretical construct that will acknowledge the role of subjective human judgment in the application of quantitative methodology to squishy problems and that will provide more explicit guidance in the exercise of that role than do

existing theories. What follows is an attempt to begin to assemble such a construct.

The application of quantitative methodology involves the use of mathematical models to represent more complex phenomena, as a way of understanding and coming to grips with those phenomena. As such, it may be viewed as a special case of a more general phenomenon--the use of models to structure and understand complex phenomena, with the term, *model*, now used more generally to denote a simplified representation of a complex phenomenon.

MODELS

The world is a complex, detailed, and confusing place. It can provide us with a far greater richness of detail about almost any problem or question that might concern us than we could possibly use efficiently. This is true at our most basic level of awareness--our continuing perception of the environment immediately surrounding us. It is also true with regard to the more complex questions and subjects that we come to know about and understand through the use of our reasoning abilities.

We cope with this plethora of information through the use of mental *models*--internalized mental structures that allow us to organize, summarize, and structure the information already available to us; to filter or screen available or incoming additional information; and to interpret and place into context that which we allow through the filter. We use our existing model of a particular situation to interpret new information about that situation, and at the same time, may use the information to modify the model. The model is not determined solely by the information available, but depends on our past history, the ways and times at which the information was acquired, and other factors. Because of the extent to which our perceptions depend on our mental models, apparently identical "objective data" may lead different people (with different models) to very different conclusions.

Even in perceiving the immediate environment, our brains develop visual perceptual models that organize and filter incoming visual information. Our eyes are broad-band sensors providing us with an

immense continuous stream of information. Relative to what we need to know about our environment to function in it, however, most of that information is redundant or superfluous or both. The load that would result from attempting to give it all equal attention all the time would be overwhelming. Our mental process organizes itself, i.e., builds an internal model, to recognize elements, patterns, regularities, etc., within the visual field that are meaningful in the environment we usually operate in and to ignore elements of the visual field that are not. The way in which this organization takes place--the perceptual elements identified as meaningful--depends on the individual's background and environment.

As an urban dweller, for example, my environment is rich in man-made structures and the geometrical regularities that such structures possess. My visual perceptual model considers regular geometrical patterns--lines, corners, squares, etc.--as significant perceptual elements. These elements occur far less frequently in nature, so presumably would play a significantly smaller role in the perceptual model of, say, an Australian bushman. At the same time, his perceptual model recognizes features of his environment that allow him to track prey, find water, and perform other tasks necessary to live in his environment. I don't know what these elements are. In fact, there may not even be names for them in English. They would likely be as meaningless to me as the apparent (to me) regularities of my environment would be to him. As a result, I might be totally oblivious to an animal trail that was obvious to him. He might be similarly incapable of picking a (to me, obvious) street sign out of the welter of an urban street scene.

To take another example, when I look at English language script, it resolves itself immediately into meaningful letters, words, and ideas (if the writer has readable hand-writing). If I look at Russian or Arabic script on the other hand, not even the letters resolve themselves immediately into meaningful units.

Similar structuring occurs in problems requiring the more explicit reasoning. "If I shoot an artillery shell into the air, where will it come down?" for example, or "What military forces should the U.S.

maintain in order to deter general nuclear war?" The amount of data available that may potentially bear on either problem is enormous. For the ballistics question, it might include the weight of the projectile, the company that manufactured it, the azimuth and elevation of the gun, the state of the gunner's digestion, the time of day, the size of the powder charge, the phase of the moon, etc. In the case of the deterrence question, it might include technical characteristics of military forces of the United States, Soviet Union, and other nations; the intelligence collection and analysis systems that provide each nation with information about the others; long- and short-term political/military objectives of both sides; the personal dispositions of present leaders and potential future leaders; the views each has of the other; etc. To answer the question, some of this data must be identified as useful, relevant information, and most must be discarded as irrelevant to the question at hand. That which is considered relevant must be structured, related, and evaluated. Models of the relevant phenomenon, e.g., ballistics or the process of conflict between nations, provide the framework within which this occurs. Depending on the nature of the problem and the person considering it, the model may be formal, precise, and explicit, or informal, vague, and implicit. It may be internally consistent or fraught with contradictions.

The two examples cited above, ballistics and deterrence, illustrate an important distinction in the types of problems we ask our reasoning capabilities to deal with. The first--the ballistics problem--is a solid problem, in the sense that a significant consensus exists among people who believe they understand the problem about what the relevant data are and how they should be used to answer the question. Most people would agree that projectile weight, azimuth and elevation of the gun, and powder charge are significant variables, and the others listed above are irrelevant. People who feel they understand the question agree on the model that should be used to answer it. This, in turn, assures agreement on the answer. The deterrence problem, on the other hand, is considerably more squishy. Among people who feel they understand the problem, considerable disagreement might exist about which elements are relevant and how they are related.

Different people, faced with the same objective data, might choose to employ different portions of that data in different ways, and could reach considerably different conclusions. Even a single individual might find himself vacillating between alternative models, with no clear cut basis on which to make a choice between them.

A model is, in effect, an economizing device, reducing the limitless number of possibilities (if considered in all possible detail) to a far more limited number of categories, i.e., the distinct states identified by the model. It does this by ignoring much of the information potentially available, by identifying or equating particular variables as always occurring in specified combination. It allows me to think of the world in terms of these categories and greatly simplifies the perceptual and processing problems I face in dealing with that world. The usefulness or lack thereof of a model is highly context dependent. It depends on how well the parts of the environment with which the models attempt to deal are represented by the categories recognized by the model, and on the nature of the penalties incurred for whatever deficiencies exist.

A model that discriminates very accurately at one place or point in time may do very poorly at another. My internal model for classifying people by sex a decade ago, say, nearly automatically identified "shoulder length hair" with "female." This identification was almost always correct, then. Identification of these same variables today would produce a high error rate. The model that I use for crossing the street tells me that, if I look to the left as I step off the curb and see no traffic, then it is safe to proceed to the middle of the street before looking the other way. In England or Japan, where traffic drives on the left, that model would be highly dysfunctional.

The extent to which the categories in our models determine our perceptions of "reality" is far greater than we usually realize. Long experience, for example, has taught us that there are four suits (categories) of cards--hearts, diamonds, clubs, and spades--and that the former two are red and the latter two black. Experiments have been performed in which subjects are shown and asked to identify a series of playing cards flashed before them from a deck containing

an occasional red spade.* If the cards are presented just slowly enough that the subject can identify them, he will identify, say, a red six of spades as a normal six of spades or as a six of hearts without hesitation. In other words, he will simply place it into an existing category that it seems to match. Slow down the speed of presentation somewhat, and he will continue to identify it that way, but he will begin to feel uncomfortable about the identification. Slow it down still further, and he will eventually notice what is wrong and identify the card correctly, most of the time. Once he has identified a few trick cards correctly, the average subject can identify additional red spades much more quickly, because he now has a category for "red spades" in his model. A few subjects fail to identify the trick cards even after repeated exposure at long exposure times. The experimenters attribute this to a tendency to "fixate after receiving a minimum of confirmation."

Similar phenomena can occur when conscious reasoning is involved, no matter how apparently careful the reasoning. If the model being used does not contain the right category to allow recognition of an event that is occurring, that event may go unnoticed, no matter how strong the evidence for it. Such is the stuff of which intelligence failures are made.

A particular model will place a new perception or piece of information in the model category in which it best seems to fit. The way I categorize a new perception thus depends on the model I use to categorize it with, i.e., on the context into which I place it. This phenomenon is illustrated in Fig. 6. Read down, and the middle *number* is 13, read across and the middle *letter* is B. As with earlier examples, similar phenomena can occur in more complex reasoning and interpretative processes. This explains, in part, the fact that the same information can sometimes be used to support, or even "prove," conflicting conclusions.

*Jerome S. Brunner and Leo Postman, "On the Perception of Incongruity: A Paradigm," *Journal of Personality*, Vol. 18, 1949, pp. 207-223.

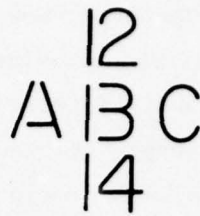


Fig. 6 — Context counts

Viewed through, or in the context of, a particular model, distinctions that are ignored or unreflected in the model may appear superfluous. Thinking in terms of ordinary cards, for example, it is superfluous to specify that a heart is red or a spade is black. Thinking in terms of hair styles of a decade ago, it seems almost a foregone conclusion that someone with shoulder length hair would be female. To examine the model, however, and to reevaluate its structure, requires that such distinctions be made. Some of the distinctions made here may appear redundant or superfluous, but I believe they are relevant to the task in which I am engaged.

Models may be classified, i.e., the set of models may be modeled, in many ways, with the utility a particular classification depending on the context. I will characterize models in different ways as I proceed, depending on the context at the time. At this time, I want to consider two characterizations: internal versus external, and logical versus intuitive.

Internal versus External Models

The *internal* versus *external* distinction is one of where the model sits relative to the perceiver's head. *Internal* models are the ones I carry around in my head, while *external* models are those I write down, draw, put into a computer, etc. This is a distinction that may appear superfluous at first glance. It becomes important, however, when we

inquire into the meaning of phrases such as "but everybody understands that" or "people who work in the area know what that means," when these phrases are used to justify models, conclusions, etc., which, taken very literally, would be obviously incorrect. The ultimate basis for an individual's personal knowledge resides in his internal models, in some very real, though hard to make precise, sense. This is not inconsistent with the fact that he may sometimes choose to use an answer provided by an external model, e.g., to a ballistics trajectory question, in preference to any he could produce internally without the use of external aids. External models can be defined by language--mathematical equations as well as words--or can be defined in other ways. Scale models, photographs, perspective drawings, etc., are also examples of external models. The corresponding things that they evoke in my head are internal models, as are the internal constructs that let me drive a car, recognize my friends, decide that someone is angry at me, etc.

Logical versus Intuitive Models

Related to the internal/external distinction, but different from it, is the distinction between *logical* and *intuitive* models. Logical models can be given an external structure, while intuitive models are largely internal. It is also possible to have an intuitive understanding of logical models. It is this type of understanding, for example, that allows good mathematicians to do good mathematics.

A *logical* model is a model whose structure is defined by an explicit set of axioms, premises, assumptions, etc. These premises, together, perhaps, with additional particularizing data such as input parameters, determine the remaining structure and behavior of the model. The model thus produced is explicit and well defined, and its behavior can be deduced as logical consequences of its defining premises. Given a question about the model, it should be possible to determine whether or not it has a well-defined answer, and if it does, to answer it, from the logical structure of the model alone, without recourse to outside information or judgment. Different people looking at the same question should arrive at the same answer on the basis of logic alone.

Arithmetic on the real numbers, for example, is a logical model, as are Newtonian mechanics and Euclidean geometry.

Logical models play a significant role in mathematics, the physical sciences, and systematic analysis generally. Because conclusions can be drawn from the model as logical consequences of its structure, agreement on the model and its inputs is tantamount to agreement on its conclusions. The explicit nature of the model's structure not only provides a basis for understanding the phenomenon being modeled, but also provides a basis for understanding the way that phenomenon is understood, i.e., for understanding the model itself. Logical models have been employed with a high degree of success in the physical sciences and engineering--areas of endeavor that contemporary society holds in high esteem. Part of this esteem has rubbed off on the models themselves, as well as on the methods of analysis that appear to use such models and, thus, appear scientific.

Logical models have a clearly defined inside and outside--with a sharp boundary between. The premises of the model and the logical consequences of those premises define the structure of the model--the inside--and everything else is outside. The edge between inside and outside is, thus, sharply drawn. One requirement imposed on the premises defining a logical model is that they be self consistent--that the contradiction of any one not be derivable from the others. This insures an internal consistency to the structure of the model as a whole.

Intuitive models, on the other hand, may not have well-defined boundaries, and frequently are not internally consistent. They are, instead, loosely defined, encompassing different parts of our overall knowledge as we shift our thought pattern from one aspect of the phenomenon being considered to another. Intuitive models consist of a number of different layers, piled, in effect, one on top of another, with each layer providing the supporting substructure for those above. We normally operate as near the surface of the overall structure as possible, probing into deeper layers only when issues arise that cannot be resolved at the surface, or when we wish to reexamine the foundations of our knowledge more thoroughly than usual. We may strive

for consistency in the part of the model that we use directly, but do not always have consistency throughout the model as a whole. When we find an inconsistency, we may choose to resolve it by carefully examining and changing some of the premises that lead to it, or we may leave it unresolved and merely suppress it by shifting attention away from the inconsistency to areas that we can deal with more easily and comfortably.

Corresponding to these two types of models, we can distinguish between logical and intuitive thought. Logical thought involves a step-by-step progression from premises to conclusions, with each step following logically from what comes before. Intuitive thought, on the other hand, may involve considerable jumping and skipping about, looking for new insights and building a holistic gestalt. Logical thought provides a means of systematically evaluating evidence within a well-defined logical structure. Intuitive thought allows the (sometimes subconscious) consideration and integration of diverse evidence, including that which may seem logically unrelated to the question at hand. It provides the mechanism for recognizing inadequacies in and changing the existing logical structure when required. These two forms of thought are complementary rather than competitive. Thinking logically about our intuitive models is a way of making them precise and verifying the insights they provide, while thinking intuitively about our logical models is the way we understand and gain new insights about them.

SURROGATES AND PERSPECTIVES

Once a model has been built, it has a structure of its own that is independent of that of the problem or phenomenon it was constructed to represent. This structure may be physical in the case of a scale model, mathematical in the case of a mathematical model described by systems of equations, verbal in the case of a model described by natural language, etc. The existence of this structure makes meaningful the distinctions between the model and the substantive problem or phenomenon being modeled, between the logic of the model and the logic of the substantive problem, between truth in the model and the truth

in the substantive problem, etc. In effect, this structure defines a problem, process, etc., in its own right that is related to the substantive problem being modeled, but is distinct from it. This structure has been referred to as "the problem defined by the model."

The distinction between the substantive problem and the problem defined by the model is important in distinguishing between two ways in which an individual may use a particular model in problem solving. He may accept the structure of the model as a valid representation of the structure of the substantive problem, and use the model as a *surrogate* for the problem in the inquiry. In using a model as a surrogate, in effect, he adopts the problem defined by the model as the problem he wishes to solve. On the other hand, he may incorporate additional information, knowledge, etc., about the problem that is not reflected in the model into his use of the model, using the model as a *perspective* on the problem.

The idea of using a model as a *surrogate* for a problem is a familiar one--so familiar, in fact, that we may not think about the fact that it is only one way, and there may be others, until that fact is explicitly pointed out. The ideas of "model" and "surrogate for problem" are thoroughly wedded in our formal education system--as wedded, say, as "shoulder length hair" was to "female" a decade ago or as "black" is to "spades" when we think about playing cards. The way to use models, we are taught (particularly mathematical models), is to build a model that captures the relevant features of the substantive problem, then solve the problem defined by the model. Then accept that solution as applicable to the substantive problem. In doing so, we too often equate the model with the substantive problem and neglect the fact that there is a very real difference between the two.

Newtonian mechanics, for example, exists as an abstract mathematical system independent of any correspondence to observed reality. It also happens to fit observed reality well in most situations. When we use Newtonian equations for trajectory calculations, we are computing the behavior of an abstract mass subjected to abstract forces in an abstract Newtonian universe, and using that model-defined problem as a surrogate for a substantive problem about real masses and forces.

The same is true of the use, say, of Euclidian geometry in surveying or even of the use of arithmetic on the real numbers in everyday counting, computing, etc. That we can ignore the problem/model distinction is a consequence of how well the models fit observed reality (like "black" and "spade"). When the fit is less perfect, however, recognition of the distinction can become important.

The use of models as *perspectives* is also familiar. We do it all the time. We usually do it subconsciously, however, so that we don't think of it explicitly as a way of using models, distinct from their use as surrogates. The nature of this use of models is best illustrated by the application from which the term is derived--two dimensional perspective drawing of three dimensional space.

The drawing in Fig. 7 is a two dimensional model of a three dimensional cube. As such, it has a two dimensional structure of its own, independent of the existence of such things as cubes and three dimensional space. Now consider the question, "What is the shortest distance from a to b on the surface of the cube?" Using the model as the surrogate for the cube, the way to answer that question is to draw the

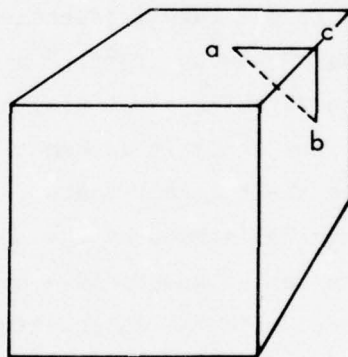


Fig. 7 — A model as a perspective

shortest distance from a to b *on the model*, as shown by the dashed line. When confronted with that problem, however, few people would choose that answer. Most would select the solid lines connecting a to c and c to b. In so doing, they are bringing to bear more knowledge about *cubes*, and distances on the surfaces of cubes, than is reflected in the

two dimensional model. They are incorporating this knowledge in their use of the model to answer the question.

In effect, the individual who chooses the lines acb as the shortest distance between a and b is using the drawing to help trigger a richer and more complex internal model of the cube. He uses a combination of the external model provided by the drawing and the internal model that drawing triggers to answer the question. In so doing, he avoids the wrong answer he would be led to by staying strictly within the two dimensional problem defined by the model.

Most of us have reasonably well developed internal models of three dimensional space, and of the way that three dimensional space is represented in two dimensional perspective drawings. This allows us to make the interpretation necessary to see three dimensions where only two exist with no conscious effort. Someone who has grown up in a primitive society, however, without exposure to perspective drawing, may be unable to readily make this transformation. There is a definite interpretive act going on. A use is being made of the two dimensional model that requires more knowledge than is provided in that model.

Prisoner's Dilemma as a Perspective

This use of a model to gain a perspective on a larger problem is not unique to the use of two dimensional representations of three dimensional objects. It is just somewhat easier to see it occurring there. The phenomenon, however, or something like it, occurs whenever any model (or a conclusion reached using one) is interpreted in a larger context than the model itself represents. Consider the game "Prisoner's Dilemma," with the payoff matrix shown in Fig. 8. The left-hand entry in each square represents the preference of Player 1, the row player, while the right-hand entry represents the preference of Player 2, the column player. The entries reflect ordinal preference only, i.e., the order in which outcomes are preferred, but not necessarily by how much.

Action C for either player should be thought of as "cooperative" behavior, while action N is "noncooperative," in the sense that it increases the payoff to the player choosing it at the expense of the

		Player 2	
		C	N
Player 1	C	(3,3)	(1,4)
	N	(4,1)	(2,2)

Fig. 8 — Prisoner's dilemma

other. Notice that whichever action Player 2 chooses, Player 1 is better off if he plays noncooperatively. Hence, noncooperative behavior is clearly optimal for Player 1. The same argument holds for Player 2, hence noncooperative behavior appears to be optimal for him as well. The choice of noncooperative behavior by both, however, leads to the (2,2) outcome, rather than the (3,3) outcome obtainable if both play cooperatively. The combination of their individually "rational" choices is collectively irrational, in the sense that it leads to a collectively suboptimal outcome.

The name "Prisoner's Dilemma" derives from imbedding the game in the following scenario. You and I rob a bank together and are caught under circumstances providing the state with a very weak case against us. If we both remain silent, the evidence against us will not support a robbery conviction. We are likely to get off on a lesser offense, such as carrying a concealed weapon. We will probably draw six months each in jail. The prosecutor would like very much to obtain a felony conviction, even at the expense of letting one of us go. He offers you the chance of turning state's evidence at my expense. If you confess and I do not, he will let you off scot free while sticking me with the maximum sentence of, say, ten years. You also know that his deputy is making the same offer to me. If we both choose to confess, he can't let both of us off. He will prosecute both of us on the felony, but recommend a reduced sentence of, say, five years each.

Should you choose cooperative (no confession) or noncooperative (confession) behavior?*

* Notice that "cooperative" here refers to cooperation between us players, and not to cooperation with the prosecutor.

confession will reduce your sentence from 6 months to nothing, while if I do confess, your confession will reduce your sentence from ten years to five. Hence, confession looks like the better choice for you, whatever you do. The same reasoning applies to my choice. Whatever you do, I am better off to confess. This reasoning would lead both of us to confess, resulting in a five year prison term for each, when we could have gotten off with six months by cooperating with each other.

Loosely interpreted, Prisoner's Dilemma provides an integrating perspective on a dilemma that occurs in varying ways throughout human social interaction, as the following examples illustrate:

- o Whether everyone else pays his fair share of taxes or not, it is always to my advantage, if I can get away with it, not to pay mine. But if everybody avoids taxes, organized society falls apart.
- o Whatever other cities along a river do, it is to the advantage of each city individually to dump raw sewage into the river rather than to construct and maintain an expensive treatment plant. If everyone does this, however, the quality of the river deteriorates for all.
- o If shepherds graze their flocks on common grazing land, each shepherd individually can enrich himself by adding a sheep or two to his flock. If all do this, however, it may result in overgrazing and economic disaster.
- o Each nation in a defensive alliance may find it individually advantageous to skimp on its contribution to the alliance. If all adopt this strategy, however, the alliance may become too weak to function.

The common theme throughout these examples is that of a collectively undesirable outcome resulting from a decision by each of the parties to attempt to maximize his individual good.

While the game of Prisoner's Dilemma provides a useful integrating perspective on the underlying dilemma common to all of these problems,

it provides very little insight into the feasibility or desirability of various ways of coping with that dilemma in different situations. The reason for this is that the solution to a particular "Prisoner's Dilemma" is far more context dependent than the dilemma itself. Depending on the particular context, the solution may contain elements of coercion, criminal sanction, mutual trust, and living with degraded outcomes. The last, in particular, tends to occur when none of the others are employed--as has been the case with various types of environmental degradation. The value of Prisoner's Dilemma as an integrating perspective on these situations is not degraded by its inability to prescribe solutions, however, since its value lies precisely in the fact that it captures the core elements of the dilemma common to all, and those elements alone are insufficient to prescribe a way out.

There is a great deal of research effort devoted to finding general "solutions" to Prisoner's Dilemma in some expanded mathematical context. This effort includes activities such as having psychology undergraduates play long strings of Prisoner's Dilemma games, with or without side payment such as the ability to apply electric shocks to their "partners;" and theoretical extensions of the game to include communication, bargaining, and side payments of various types. Some researchers engaged in this type of activity seem to feel that Prisoner's Dilemma, as it stands, is an anomaly that must be done away with, a paradox to which there must be a "rational" solution. They find the existence of the game in its pure form unacceptable. I believe to the contrary, that the pure form of the game is a useful perspective on a common human dilemma because of its stark simplicity, because it does capture the central elements of that dilemma in a context free way. Attempts to extend the applicability of the model further--by looking for solution within the structure of the model itself or by embodying it in a larger mathematical context--are as likely to detract from its perspective value, in my opinion, as to add to it.

The Pueblo Risk Assessment Process

Another example of the use of a mathematical model as a perspective can be found in my study of the risk assessment process for the

mission on which the U.S.S. Pueblo was seized in January 1968.* In that study, a decision theoretic model of the risk assessment problem for such missions was constructed and used as a source of standards against which the risk assessment process for the Pueblo could be measured. No assumption was made that complex operational risk assessment problems were amenable to quantitative treatment. In fact, the converse was argued. Nonetheless, the model was found useful as a perspective (although that term was not employed) on the risk assessment problem from which useful observations about risk assessment could be drawn. The philosophy underlying this use of a mathematical model was summed up as follows:

In order to evaluate the risk assessment process for the Pueblo mission, some standards against which to measure that process are required. In order to obtain some standards, we will turn to a mathematical discipline in which problems of risk assessment are rigorously addressed--statistical decision theory. That theory deals with and provides quantitative tools for risk assessment in abstract mathematical models or in problems which possess a sufficiently well-defined and quantifiable structure to be adequately represented by such models. Unfortunately risk assessment problems such as the one represented by the Pueblo mission are not in that class. Hence, the tools provided by statistical decision theory are not directly applicable.

The theory, nonetheless, has value in dealing with such problems. The basic structuring of the problem suggested by the theory, and the conceptual principles used, are valid regardless of whether or not the detailed structure of the problem as such is quantifiable. These principles, then, interpreted qualitatively and stripped of their quantitative content, may act as guidelines for risk assessment when quantification is not possible, and as standards against which to measure the risk assessment process in unquantifiable situations such as that represented by the Pueblo mission.†

* *The Operational Assessment of Risk: A Case Study of the Pueblo Mission*, The Rand Corporation, R-691-PR, March 1971.

† *Ibid*, p. 21.

Perspectives and Surrogates in Policy Analysis

The perspective/surrogate distinction is an important one in the application of quantitative methodology to policy analysis. Many questions occurring in policy analysis are sufficiently squishy to preclude construction of an adequate surrogate, in the sense that models in the physical sciences are surrogates. Rather, the fit between model and problem in most squishy problems is such that any model is best viewed as a perspective--one way, but not a complete or uniquely valid way, of viewing the problem. In most situations, the perspective provided is likely to be sufficiently incomplete to require considerable use of additional substantive knowledge to determine how the model should be used and how the results obtained should be interpreted.

But the characteristic of being a perspective or surrogate is not one that resides in the model alone, or even in the combination of model and problem. It resides, rather, in the head of the analyst using the model--in the way he thinks about the relationship of model and problem in using one to understand the other. Quantitative methodology and its supporting theory have been developed with an eye toward employing models as surrogates rather than perspectives (in part by default, since the distinction is not explicitly made). When this methodology is applied to squishy problems in a method-oriented way, we frequently end up with what might be called "pseudo-surrogates." These are models that are clearly inadequate as surrogates for the problems to which they are applied, but that are treated as though they were (with appropriate caveats) in the hopes that the users of analysis, or "decisionmakers," can interpret them in a broader substantive context, i.e., can make use of them as perspectives. How well people do this is subject to question, and since it is something that goes on internally, it is not a question subject to objective investigation. It seems likely that the pretense involved detracts from the quality of understanding that could be achieved through more open acknowledgment that the models are, at best, perspectives, and by approaches to analyses that take explicit account of that fact.

One of the major differences between a perspective on a problem, and an adequate surrogate, lies in the relative nonuniqueness of the

former. Loosely speaking, we expect an adequate surrogate for a problem to be unique and to provide a unique answer or solution to the problem.* A perspective, however, is nonunique in two important ways, both of which are illustrated in Fig. 9. First, there is no single valid perspective on a given problem. It may, in fact, look quite different from different perspectives. Second, depending on the perspective chosen, very different problems may look the same. Both forms of non-uniqueness have significant implications for the use of perspectives in policy analysis.

The first form of nonuniqueness--the existence of different perspectives on the same problem--is important because of its implications for the resolution of apparently conflicting or inconsistent views. If two models yielding different answers to the same question are both viewed as surrogates, then it seems reasonable to assume that one of them must be wrong and a choice must be made between them. If they are viewed as perspectives, however, the choice is not so simple. It may be that both are valid, and that the proper resolution of the conflict between them lies not in accepting one and rejecting the other, but in coming to understand the problem in a way that admits the validity of both. A classical illustration of this point is the story of the blind men examining the elephant. One felt the trunk, and described the elephant as being like a snake; another felt a leg, and described it as like a tree trunk; a third felt an ear, and described it as like

* Strictly speaking, of course, there is no unique surrogate either. Newtonian mechanics and relativistic mechanics, for example, are both useful surrogates for "real mechanics." Similarly, either the probability space consisting of two points, each of which is assigned probability one-half, or a probability space supporting an infinite sequence of independent, identically distributed, random variables, each 0 or 1 with probability of one-half, with one of those variables identified as the one of interest, can be used as a surrogate for one toss of a fair coin. Generally speaking, however, when two valid surrogates with the same problem exist, one of them may be viewed as an approximation of the other (as in the first example), or as a subsystem of or surrogate for some portion of the other (as in the second). Both would be expected to provide approximately the same solution, and if they do not, the question of which one is right is a meaningful one. This need not be true in the case of perspectives.

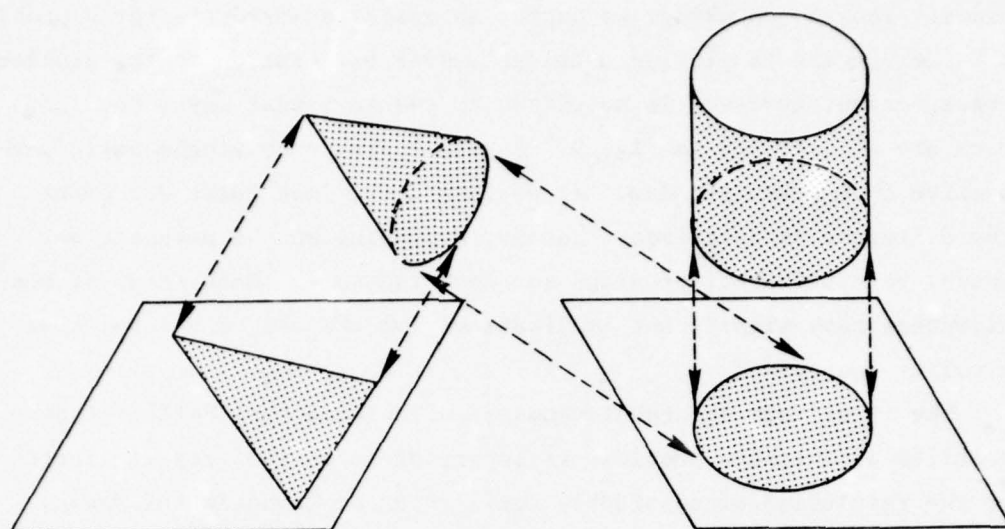


Fig. 9 — Nonuniqueness of perspective

a leaf; and the fourth felt its side, and described it as like a wall. If we feel compelled to accept one man's description and reject the others, then these apparently inconsistent descriptions are clearly irreconcilable. Given our fuller understanding of what an elephant is, however, it is easy to see that they are reconcilable as distinct perspectives on a complex object. All too often in policy analysis, arguments develop over which of apparently conflicting analyses is "right," and which is "wrong," when neither is, in fact, completely valid or invalid. More open and explicit acknowledgment that perspectives are being used might frequently suggest instead that attempts be made to combine the two, and see what insight both provide beyond that available from either one alone.

The other form of nonuniqueness--of the problems that a given perspective represents--has implications for the interpretation as substantive conclusions of the results obtained by analyzing a perspective. On the one hand, the results may be such that the corresponding conclusion holds for every problem for which the perspective is valid. On the other, they may be such that the validity of the translation from results in perspective to substantive conclusions about the

problem is highly problem dependent--true for some problems for which that model is a valid perspective and false for others. In this latter case, it is incumbent on the analyst to insure that his problem is one for which the translation is valid, and not simply to make that assumption without question.

In verifying the gross quantitative sufficiency of strategic forces for the deterrence of nuclear war, for example, the "gnome in the basement" model of Soviet decisionmaking is an adequate perspective on a broad range of possible "real" Soviet decisionmakers and decisionmaking mechanisms. Strategic forces sufficient to deter the "gnome in the basement" are likely to be sufficient to deter any Soviet decision-making mechanism concerned with insuring the survival of the Soviet Union as a viable society--however that decisionmaking mechanism evaluates viability, "unacceptable damage," etc. The conclusion that forces capable of inflicting AD on the Soviet Union will deter actions that the Soviet leadership perceives as triggers to that response, then, is not highly sensitive to the question of what "true" Soviet decision-making mechanism the "gnome in the basement" happens to be a perspective on.

This same perspective--the gain/loss calculating gnome--may superficially appear useful for inferences about other forms of Soviet behavior. He has been used, for example, to support the conclusion that the Soviet Union would be strongly motivated toward a counterforce attack on the Minuteman force by the fact that *our* force exchange analyses show the force to be potentially vulnerable. Even granting the validity of the engineering calculations making up such analyses (which may themselves be somewhat dubious), their interpretation in political motivational terms is highly questionable. Such conclusions follow from the same inherent logic *in perspective* as does the AD calculation, but are applicable to a far smaller set of problems on which that perspective might be drawn. Instead of retaining their validity across a broad range of Soviet decision mechanisms (those for which "real" Soviet leaders positively value the survival of their society) as does the AD calculation, they require that the Soviet leadership make essentially the same calculations and see the problem in essentially the same terms as does the "gnome in the basement" used as the perspective.

One of the potentially more serious difficulties created by the nonuniqueness of the problem represented by a perspective occurs because problems change with the passage of time. I may create a model at one time that is an adequate and useful perspective on my problem at that time. Over time, however, the problem may change in ways that invalidate the perspective but that are not themselves reflected in or through that perspective. Thus, if I view the problem only through that perspective, I may never notice the change taking place. My initially adequate and useful perspective may gradually become inadequate and distorted in ways I never notice. A simple example might be the shifting line, in terms of goods or purchasing power, between petty and grand larceny resulting from the combination of inflation and a constant dollar definition.

Analysis using a model as a perspective is both more and less demanding than analysis using a surrogate. It is less demanding in technical model building terms, because it places less value on a high degree of "fit" between problem and model. The model need not include everything that "counts" in some abstract sense, but only those aspects of the problem that are required for the perspective the analyst wishes to investigate. This is because he will not restrict himself to analysis of the problem defined by the model, but will bring in and use his additional knowledge about the problem throughout the course of the analysis.

This same factor (looked at from the other side) is what makes analysis using a perspective more difficult. If the model is acknowledged as an inadequate surrogate, the analyst cannot simply retreat within the model and justify results on the grounds that they follow from the logic of the model. He cannot depend to the same degree on rote compliance with rules of procedure, and on the validity of his logic within the model, as grounds for his substantive conclusions. The shortest distance between two points on the surface of a three-dimensional object may be very different than the apparent shortest distance between the same two points on a picture or a perspective drawing of that object. For this reason, it is incumbent on the analyst to remember that he is dealing with a perspective and not to attempt to treat it as a surrogate.

He cannot, for example, uncritically rearrange the elements of his perspective in ways that seem logically permissible *within the model* and be sure that what he will get will be valid. If he attempts to do so, he may well produce something whose elements each look reasonable in perspective, but fit together in a way that makes them absurd as a representation of reality. Both the left and right halves of the object in Fig. 10, for example, are valid perspectives on objects that exist in three dimensions. The overall drawing is not.

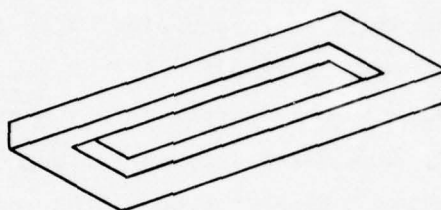


Fig. 10 — A risk of working in perspective

This same phenomenon occurs in more complex problems and models, as, for example, in strategic analyses that seem to conclude that political leaders would embark on global strategic conflict in order to win "points" according to a scoring system that analysts find convenient for the evaluation of weapon systems. The analyst lays out a series of assumptions (each of which seems plausible as a perspective on reality), analyzes a model embodying those assumptions, and interprets the results thus produced as substantive conclusions. He justifies his interpretation on the grounds that "the conclusions follow from the assumptions."

If the individual assumptions are "surrogates" having the validity of, say, the law of physics, this line of argument is valid. If they are approximations or "perspectives," however, it is not. Evaluation of the assumptions individually must be supplemented by holistic evaluation of the structure produced by the total set of assumptions. The argument that "The assumptions are acceptable individually, and the conclusions follow from the assumptions; therefore, the conclusion must be acceptable," is simply a parallel to the following argument

about Fig. 10. "Something can be built (in three dimensions) that looks like each half individually; therefore, it must be possible to build the whole thing."

What's happening here, in the argument about Fig. 10. and in more complex cases where "the conclusions follow from the assumptions," is that we are being asked to make judgments only about the assumptions individually, and then to suspend judgment and accept whatever model and conclusions follow from those assumptions. When the model being constructed is taken as a surrogate, as in the case, say, of a ballistics model, this procedure is reasonable. When the model is at best an imperfect perspective, it is not. If the results obtained from the model are to provide any real insight into the problem, the continued exercise of careful, considered judgment throughout the remainder of the analysis, including interpretation of the results, is required. Whether this is done or not is not a property of the model, but of the analyst using the model and of the way he uses it. The conclusions thus produced are the *analyst's* conclusions, albeit based on the model, *not* the model's conclusions.

The analyst, therefore, bears, whether he likes it or not, a degree of personal responsibility for the conclusions he produces. He cannot abdicate that responsibility by suspending sentient judgment following the choice of his assumptions, grinding on to some of the logical implications of those assumptions, then arguing simply that "the conclusions follow from the assumptions."

This is not, of course, what happens in good quantitative analysis, and there is a great deal of good quantitative analysis. It is, however, what happens in much bad quantitative analysis, and there is, unfortunately, a great deal of that also. This approach, of suspending judgment following the choice of the assumptions defining the model and accepting as valid the logical implications of those assumptions, is inherent in the formal theory from which quantitative methodology is drawn. We thus find ourselves in the unfortunate position that major abuses of the methodology are strongly related to literal interpretations of the supporting theory.

Quantitative methods have significant potential as a policy analysis tool, when used in conjunction with carefully considered human

judgment. This requires that the models employed be thought of, and used, as perspectives rather than as surrogates. Good analysts do this. Unfortunately, this use of models, in conjunction with additional knowledge and judgment outside the model on the part of the analyst, is not dealt with by the supporting theory nor adequately recognized in the conventions that exist concerning the use of quantitative methods. We seem to be in a situation in which the extant interpretations of the methodological theory and the conventions concerning its application are major contributors to the abuse of the tools provided by the theory. This suggests a strong need for a serious rethinking of those interpretations and conventions.

RIGHT AND LEFT THINKING

It could be argued that the surrogate/perspective distinction is not a meaningful one--that models are never really full surrogates for the problems they represent, but rather approximations to those problems. The question is, then, one of how closely the model approximates the problem, with different degrees of closeness being differences in degree, but not in kind. This argument misses the essential point of the surrogate/perspective distinction, which relates not to the question of how closely the model fits the problem objectively, but rather to the question of how the analyst uses the model in combination with additional knowledge he possesses that is not incorporated in the model. It is a question of how the analyst thinks.

Now thinking is something we all do, at least occasionally. Thinking about the way we think, however, is something we do less frequently. Thinking about how the way we think should affect our use of and our expectations about tools that aid our thought (such as computers and quantitative methodology) is something we do only rarely. This last question goes to the core of the issue of the judgemental application of quantitative methodology.

Man has long recognized the existence of two types of thought, described as logical versus intuitive, linear versus holistic, intellectual versus sensual, verbal versus nonverbal, etc. The verbal/logical process is often thought of as dominant among scientists,

lawyers, mathematicians, etc., and the nonverbal/intuitive process as dominant among artists, poets, and mystics. Since contemporary Western society values the former skills more highly than the latter, we tend to value the verbal/logical thought process more highly than the nonverbal/intuitive process. This valuation is reflected in many ways, particularly in the education process and in the emphasis that contemporary education places on verbal and logical skills. This has the effect of denigrating the value of the nonverbal/intuitive process and training people to avoid and mistrust the process.

A growing body of evidence is accumulating to support the hypothesis that two separate thinking processes do exist in the human mind, and in fact take place in different sites in the brain.* The verbal/logical/linear process takes place predominately in the left hemisphere, while the nonverbal/intuitive/holistic process takes place primarily in the right. Evidence supporting this physiological separation of function comes from testing and observation of subjects who have had the hemispheres of their brain surgically separated.

Our brain consists of a right and left hemisphere. Each hemisphere receives sensory information from and controls the motor activity of the opposite side of the body. Thus, the left hemisphere receives what the right visual field sees and controls the movements of the right side of the body, while the right hemisphere performs comparable functions on the left side. The two hemispheres are physically and neurologically connected, allowing transfer of information from one to the other. Information received by a person with a normal brain is thus available in both hemispheres.

The two hemispheres are sometimes surgically separated as a treatment for epilepsy, in order to prevent a seizure occurring in one hemisphere from spreading to the other. When this is done, the two hemispheres are no longer able to communicate with each other internally. Information known to one hemisphere will not be available to the other unless it is communicated externally. In normal day-to-day functioning,

* My description of this evidence is taken largely from Robert E. Ornstein, *The Psychology of Consciousness*, Viking Press, New York, 1972.

this creates no particular problem for the split brain subject, since the major information channels feeding the brain, the eyes, and ears are parallel channels which feed both hemispheres simultaneously. The split brain subject normally shows no noticeable perceptual/motor deficit in his day-to-day functioning.

When the right and left input channels are isolated, however, the results can be striking. Allow a split brain subject to handle a familiar object, such as a pencil, with his right hand, but not to see it. He can identify it and describe it with ease just as a normal subject could. Give him the same object in his left hand, however, and he is unable to describe or identify it. The right hand feeds the left hemisphere of the brain which controls language. The left hand feeds the right hemisphere which has no language production ability. The left hemisphere, where the language is, literally does not know what the left hand is doing, and therefore cannot verbalize about it. The right hemisphere does possess the ability to interpret language, however, so if the subject is shown a printed list of objects, he can point to the name of the one which he has been handling.

Left hemisphere deficiencies in spatial perception can be exhibited in a similar manner. A common test of visual perceptual abilities requires the subject to reproduce a three dimensional pattern of blocks. Split brain subjects can do this readily with the left hand (which is controlled by the right hemisphere) but have considerable difficulty in attempting to do it with the right (controlled by the left hemisphere).

For me, one of the most striking findings of this type of research concerns the ability of the two hemispheres to develop external channels of communication. In one series of experiments, the researcher presented the subject's left visual field with the randomly chosen red or green light, and asked the subject to identify the light color he was seeing. The subject's initial responses were essentially random guesses, because his left hemisphere had no way of knowing what he was seeing. Soon, however, the subject developed a pattern of correcting himself if his initial guess was wrong. Having made the wrong guess, he would frown or shake his head, then change his answer to the correct one.

What seemed to be happening was that the right hemisphere heard the answer being given and knew it was wrong. Having no *internal* channel with which to correct the left hemisphere, the right hemisphere chose *external* communication--by causing a frown, headshake, etc.--to call attention to the mistake. The left hemisphere learned to recognize this signal and to correct itself accordingly.

Everyday experience suggests that similar "attention getting" mechanisms are used by the nonverbal thought process in normal persons. The familiar "doubletake," involving a sudden start or movement coincident with recognition of some important fact, or the occurrence of an involuntary gesture such as the slap on the head, accompanying a sudden insight, are probably examples of this phenomenon. These may be necessary because of the fact that contemporary society teaches us to ignore our nonverbal thought. Our consciousness is focused on the left hemisphere verbal process most of the time, and we pay very little attention to the right-hemisphere nonverbal process. When that process has something important to tell us, therefore, it must get our attention through external means--the doubletake or involuntary slap on the face. One of the points I am trying to make in this report is that the nonverbal part of the thought process has a lot to tell us. It is the source of most original or creative knowledge of all types--scientific or logical as well as artistic--and we need to learn to pay more attention to it. At the very least, we should learn not to ignore it when it has something important to say to us. This is not to say that we should turn away from logic and totally toward intuition, but simply to suggest that we recognize our intuition as a source of ideas and knowledge. We must still express those ideas verbally, and where necessary, validate or invalidate them logically. We have no chance to do that, however, if we suppress the signals and ignore our intuition all together when it tries to tell us something.

This point is illustrated by a conversation I had with a colleague a few years ago concerning some force posture analysis being done under Rand's strategic program and the validity of interpretations being made of some of the quantitative results produced by that analysis. My

position was that the analysis itself was overly simplistic and that it sufficiently distorted the situation being analyzed to make the interpretations questionable, if not completely specious. His position was that while these objections might be valid in principle, the analysis was being done in a way that was generally accepted within the strategic planning community and would produce answers "needed" by the client. Any attempt to deal seriously with my objections would render the problem analytically intractable and make it impossible to get "good solid answers." Besides, the assumptions were spelled out and the conclusions followed logically from the assumptions and that was what our job was--to draw logical conclusions from explicit assumptions. Deciding whether or not the assumptions made sense, and interpreting the conclusions in the light of whatever deficiencies there were in the assumptions, were someone else's responsibility, not ours. At one point in the conversation he said something to the effect that "All you're doing is voicing your gut feelings. You've got to learn to ignore your gut feelings and stop worrying about them. We all have gut feelings like that, but if we paid attention to them, we would never be able to get any good solid analysis done, and that's what we're paid for."

This viewpoint is, I believe, both revealing and representative. Many of those of whom this assessment is most critical--at least with respect to the way they do analysis--share with me, at an intuitive gut level, the same qualms and reservations about the work they do that led me to this inquiry. A major difference, perhaps, between them and me, is that they suppress those feelings in the interest of "getting the job done," while I have found myself unable to, and tried instead to understand and articulate them.

One of the reasons for the near-exclusive focus on the left-hemisphere verbal process in contemporary society is that that's where the language is, and language is our primary interface with the rest of the world (and often, even with ourselves). We know what others think through what they say and write and are aware of our own thoughts primarily when they are in verbal form. Much of our accumulated knowledge is verbal in nature, stored in the form of language. This almost

constant coupling of thought and language sometimes makes it appear that language is a necessary part of thought, at least of complex or abstract thought, and that nonverbal thought is necessarily inferior, unintelligent, etc. It seems highly likely, however, that this view is significantly in error and that the nonverbal thought process plays a considerable, if largely unnoticed, role in human intellectual activity.

To see that thought need not necessarily be coupled to language, consider the evolutionary background of language in man, and the possibility of a very different type of development in different circumstances. What we know as language today undoubtedly had its beginnings in the use of vocal cries, grunts, etc., to communicate simple ideas with positive survival value--"help," "here's food," "danger," etc. The use of noises to describe characteristics of the observed environment--"hill," "water," "food," probably came next, followed by the addition of modifiers and the use of longer strings of sounds tied together to produce and communicate more detailed and complex ideas. The invention of writing--the use of symbols to represent words--further expanded the potential usefulness of language as a means of conveying thought. First, it allowed thoughts to be recorded and kept track of in a nonambiguous way. In addition, writing provided a basis from which the ability to create new thoughts from old by manipulating the symbols involved could be achieved. This ability to manipulate and create thoughts by manipulating the symbols that represent those thoughts finds expression in mathematics, logic, etc. It is an ability that has been further developed and refined through the use of increasingly complex external aids, culminating in the digital computer. All of this seems to stem from the fact that the initial mechanism at our disposal with which to communicate with our fellows about our environment was a series of auditory grunts.

Do Dolphins Think without Language?

Is language the only way in which the ability to express and communicate complex ideas might evolve, or are there others that would not necessarily involve language at all? We do not have an adequate basis

to answer that question, definitively, either way. What is known about the dolphin, however, suggests the possibility of the development of a high degree of nonverbal intelligence in an environment that would produce no need for or pressure toward development of language. Bottle-nosed dolphins are sea mammals that are thought by many observers to be at least as intelligent as man. They exhibit highly complex behavior and have brains on the order of from 20 to 40 percent heavier than those of man. They evolved in an environment much different from that in which man evolved and, hence, may have evolved quite different thought processes. In particular, they may have no need for language as we know it.*

The dolphin lives in water, which at times may be dark, murky, or muddy. Even in the best of conditions, light penetrates to at most a few hundred feet. Vision, therefore, is a far more limited sense in water than in air. While the dolphin has and uses vision, he does not rely on it to the extent that we do. His primary sense is his auditory sense. This sense is not only passive as in man, but includes a highly evolved and complex system of active echo ranging. The dolphin probably obtains at least as much information about his environment auditorily as we do visually, and roughly the amount visually that we do auditorily.

The dolphin has three phonation (sound-producing) devices. Two of these are similar and are used in coarse resolution echo ranging and in communication (below 50 kHz). The two can be used together to produce stereophonation, or can be used as separate parallel communication channels. The third phonation device produces higher frequencies (25 to 150 kHz) and is used for high resolution echo ranging.

For coarse mapping of his environment, the dolphin uses what Lilly calls a "slash call"--a call whose frequency rises linearly with time. The length of the slash call (0.1 to 0.6 sec) is such that the call is still going out when the echo from nearby objects returns. Because

* My description of dolphins and their behavior is taken from *The Mind of the Dolphin*, by John Cunningham Lilly, M.D., Doubleday & Co., Inc., Garden City, New York, 1967. The speculation concerning the nature of the dolphin's thought process is my own.

of the time-frequency relationship in the slash call, the beat frequency (difference between the slash call and the echo) will be a function of the distance to the reflecting object. The stereo separation between the dolphin's ears will indicate the direction of the object, and the presence or absence of doppler shift will indicate relative movement or lack thereof. The slash call and its echo thus provide the dolphin with a very complete low resolution picture of his environment. In Lilly's words:

The internal picture which the dolphin can then create while sounding slash calls, the internal picture which he creates of his surroundings in terms of beat frequencies coming stereophonically combined from the two ears, must be a very interesting kind of picture. It is as if to us the nearby emitted a reddish light and the farther objects emitted a bluish light, with the whole spectrum in between. We might see, for example, a red patch in the distance farther away . . . a blue background downward symbolizing the bottom, a red patch up close meaning a fish nearby, and a large green object swimming between us and the bottom meaning another dolphin. This conversion of their acoustic beat frequencies into colors is one way we can visualize how their surrounds look to them. (Once again, as in the previous account earlier in this book, we must convert their "acoustic pictures" into our visual pictures, because of the differences in our brain and in our approaches to our surroundings.) . . .

Returning then to the slash call and the objects of varying distances from the dolphin, if he is remaining motionless in the water, he will hear various frequencies. If objects are receding from him, he will hear the frequencies dropping. Similarly, if he is moving through the water very fast and other objects are also moving through the water very fast with him, these echoes will have apparent beat frequencies which will remain constant, whereas objects which they are passing will have the rising-falling Doppler effects. Thus the internal picture of his surrounds becomes more complex and we must bring in change-of-frequency with respect to time as well as frequencies which merely depend upon distance.*

If he finds something interesting which he wishes to see with higher resolution, the dolphin can turn toward it and scan it with his high frequency, high resolution, echo ranging apparatus.

*Ibid., pp. 151-152, 154. Reprinted by permission of the publisher.

The auditory picture that the dolphin constructs of his environment, then, is probably as complete as (although quite different from) the visual picture we construct of ours. When it comes to telling another dolphin about somewhere he's been or something he's seen there, however, the dolphin has it all over man. The medium that the man has available for communication (language) is far narrower and more restrictive than the medium (visual sensation) that produces most of his information about his environment.* At best, therefore, a man can communicate only a small part of the idea of what it is "like" somewhere, or what he experienced there. The dolphin, on the other hand, may be able to communicate something approximating his full auditory experience. Using his stereophonation apparatus, he may be able to reproduce the echo he received from his slash call, and in this way, communicate the full sensation of "being there" to another dolphin. Again, to quote Lilly:

When a dolphin wishes to talk about an object at a given distance to another dolphin and wishes to describe how that object moved and at what velocity, he can do it merely by transmitting the proper frequency pattern in his clicks and whistles. In other words, he can converse about moving down from the surface of the sea toward the bottom, he can converse about fish of a given size at a given distance, sharks of a given size and all of these other matters, in a frequency-time-intensity domain which we would have to convert into visual images.†

In human terms, this would be the equivalent of direct transmission of eidetic imagery from one human mind to another.

We have no way of knowing at present what the dolphin's capabilities or predilections are for the creation and manipulation of abstract thought. On the one hand, it could be argued that the very narrowness of language--its inability to describe our world in the richness in which we perceive it visually--is what provides the impetus to

* Nonverbal communication--gestures, etc.--add considerably to the verbal in face-to-face communication in some situations. The basic limitation, however, that one can communicate only a small fraction of what one knows, remains.

† Ibid., p. 154. Reprinted by permission of the publisher.

abstraction necessary to the development of abstract thought. Since our experience is limited to but a single case, however (man), that argument is far from compelling. Abstract thought might well develop in the absence of language, and the greater flexibility and richness of detail which eidetic imagery (visual or auditory) can provide in comparison to language suggests that it could be a better mechanism for dealing with abstract thought as well as for describing the physical environment.

The purpose of this digression has not been to argue pro or con for the intelligence of dolphins. Rather, it has been to suggest, by plausible if not proven counter-example, that thought and language need not go together--that the former does not depend on the latter for its existence. This is true, in humans as well as dolphins. We "know" more about most of the things we know about than we can express verbally. It is important, therefore, that we recognize language for what it is--a tool for expression and communicating abstractions of our knowledge and perspectives on our nonverbal understanding--and not, except in certain specialized situations, a means of telling all we know or understand.

The relevance of all this to policy analysis lies in the fact that the process of analysis is a cognitive activity, a form of thoughtful inquiry. That activity produces two products. One of these is invisible--the added understanding, insight, etc., into the problem that the analyst gains internally. The other is the external rendering of that understanding--the briefing, written report, etc. The latter is a summation, model, etc. of the former, constructed of language. In problems that are sufficiently complex and squishy that they cannot be rendered unambiguously into language, literal interpretation of the language artifact may produce a highly inadequate representation of the analyst's nonverbal understanding. To go back one step further, if the analyst himself focuses too strongly on the verbal level and on understanding and solving his external model of the problem, he may achieve considerably less overall understanding of the problem than he could have achieved by giving it more careful attention on a non-verbal level--by thinking about it a little more. He may then transmit

this limited understanding onto his client and other users of his product in the false guise of insight into the substantive problem. He thus does considerable disservice to himself and his client.

VERBAL AND NONVERBAL THOUGHT

The foregoing suggests that human thought in general, and problem solving in particular, is a highly complex process whose complexities must be appreciated--if not fully understood--in any attempt at modeling human thought processing and problem solving in order to better understand judgmental analysis of squishy problems. What follows now is an attempt to construct such a model. It should be thought of as a perspective on the human thought process, and not as a surrogate for it. That is to say, it is intended to guide subjective thought about the process of thought, not to define that process. The discussion that follows, therefore, should be thought of as an unfolding description--as yet incomplete at all times--rather than as a model that will be complete at some point.

The model will be organized around the concept of two distinct types of thought--verbal and nonverbal--and the use of models as the conceptual structures on and within which those thought processes operate. Verbal and nonverbal thought are, as we shall see, quite different, although they blend together and the line between them is not always distinct.

Verbal Thought

Verbal thought deals with *language* and consists of meaningful symbols such as words, numbers, mathematical symbols, and linear strings of those symbols such as sentences and equations. The strings are also meaningful, with their meaning derived in part (but not necessarily completely) from the meaning of the individual symbols. Two most important examples of language are natural language and the language of mathematics. *Natural language* is that which we use to communicate with other people most of the time--ordinary English. The elements of natural language are words and sentences made up of words. The *language of mathematics* is language used to express mathematical

concepts and ideas, define mathematical models, derive results, etc., in precise ways. The language of mathematics includes natural language words used with precise meanings, as well as numbers, mathematical expressions, equations, etc. The term, *language of mathematics*, is used to distinguish between that language and the objects it describes, e.g., between the words "random variable" or "probability space" and the abstract mathematical objects that those words denote. Computer languages would also be languages within this definition. "Body language" or other forms of nonverbal communication would not.

One feature of language that plays an important role in determining its utility, as well as its limitations, is its linearity--the fact that it consists of a sequence of symbols following one another in linear order. On the plus side, this allows us to produce, manipulate, or process a few symbols at a time, something our verbal thought process and our auditory/vocal communication channels do well. At the same time, this linearity limits the usefulness of language in expressing holistic concepts--what it feels like to drive a car, for example, the unique facial characteristics that allow me to recognize a friend on sight, or how I understand a complex policy issue.

The verbal thought process operates on and with language. It constructs and manipulates symbols and strings of symbols, and performs operations on and with them. Speech, arithmetic, logical argument, quantitative comparison, etc., all fall within the purview of the verbal thought process. This process is sequential in nature, matching nicely with the linear nature of language. External extensions of human thought, such as the abacus, the adding machine, the "back of the envelope," or the large-scale digital computer, are extensions of the verbal thought process. They are manipulators of formal symbols, and they extend the human ability to keep track of and to manipulate those symbols.

Nonverbal Thought

The nonverbal thought process is a far different animal. The elements on which it operates are more diffuse and harder to define in verbal terms. It deals with concepts, spatial relationships, loosely

defined groupings of "sameness," etc. It deals with things as integral wholes rather than as strings of symbols, and its primary mode of processing is holistic rather than sequential. It is the place where intuition operates, where arational thought occurs, where emotions, feelings, and values find expression. It is in the nonverbal process that our internal models that reflect awareness of the environment around us reside and our responses to that environment are generated, even when these responses find a verbal form. The nonverbal thought process is one we are unable to duplicate externally. When we try to replicate a nonverbal function externally, as, for example, pattern recognition on a computer, we do so by attempting to translate that function into verbal/logical terms. It is then a different function from the one we perform holistically. It is perhaps for this reason that computers are less successful in pattern recognition and related tasks than in the performance of purely logical functions. Those tasks are performed by people in ways that computers are totally unable to emulate.

The distinction between logical and intuitive thought discussed earlier (pg. 50) closely parallels, although it is not, strictly speaking, identical with, the verbal/nonverbal distinction. Logical thought is necessarily verbal, since logic is a verbal construct. Intuitive thought is primarily nonverbal although probably not completely. It can certainly, at times, be put into verbal terms.

The Complementarity of Verbal and Nonverbal Thought

The verbal and the nonverbal thought processes are sometimes viewed as competitive--perhaps even as incompatible with one another. Because of the relative emphasis placed on it by contemporary society, the verbal process is sometimes thought of as a superior or major process, and the nonverbal as an inferior or minor process--one that any right thinking, logical man would do well to ignore if not to totally suppress. This characterization of the verbal and the nonverbal processes as competitive is fundamentally in error. It arises from a confusion of product and process--a failure to distinguish between "the footprints and the men who made them" in the words of the Chinese philosopher Chuang Tsu. The two processes are complementary rather than

competitive, working together in a mutually supportive way in most of the problems with which people deal. The nonverbal process supports the verbal by suggesting ideas, paths of inquiry, etc., while the verbal serves to express, structure, and validate the nonverbal.

Formal Euclidean geometry, for example, is superficially the most completely logical activity most people ever engage in. A proof consists of a sequence of statements, each following logically from explicit axioms or from the statements preceeding it. The terms used (line, point, triangle, etc.) are allowed only the meanings attributed to them by those axioms. No appeal to their more usual "real" meanings or to intuition about their "real" behavior plays any part in the formal theory. Nonetheless, intuition plays a significant role even here. It is the nonverbal understanding of relationships between lines, points, and angles that guides and directs the logical process of constructing formal proofs. We may pretend that the formal system is all that matters, but for most people that is not the case. What matters is a combination of formal logic and of intuition about relationships that the formal logic is manipulating. Conversely, the verbal process provides the means of organizing, recording, and expressing the product of the nonverbal process. The poet, for example, may be dealing primarily with nonverbal ideas, but he uses a verbal medium to express them.

Consciousness can be thought of as a spotlight that plays around on the contents of the mind, or as a lighted stage onto which the small part of that contents can be brought for viewing. We are generally conscious only of those thoughts that are directly in the spotlight. Sometimes, we are less directly conscious of nearby thoughts and ideas as well, in a manner similar to our awareness of our peripheral vision. Our nonverbal consciousness is primarily sensual--awareness of feelings, of listening to music, of taste, and of the combination of sensations and feelings associated with physical activities such as skiing or lovemaking. Some people are also conscious of ideas in a nonverbal (sometimes visual) form, and do much of their conscious thinking nonverbally, verbalizing the ideas only to make them precise or to communicate them with others.

For many people, however, consciousness of ideas is primarily verbal, taking the form of awareness of a stream of verbal thoughts

that they focus and direct, sometimes, as they think. For most people, most of the time, this verbal thought stream fills their consciousness sufficiently to preclude direct awareness of the nonverbal process. Even when this is the case, however, the nonverbal process is there, influencing and directing (and being influenced by) the verbal. The apparent preeminence of the verbal process in dealing with ideas arises because ideas find their most common expression in terms of language--the medium in which the verbal process operates.

The verbal process plays two distinct roles in human thought. The first is as a thought process in its own right--as a manipulator and transmitter of ideas in the form of language. The second is as an interface between the nonverbal process and the outside world, as a way of expressing thoughts generated nonverbally. This suggests an image of nonverbal thought as being, in some sense, underneath or beyond the verbal, accessible primarily through the verbal, rather than directly. It also suggests an important distinction between verbal thought at the purely verbal/logical level and verbal thought as directed by, as an expression of, etc., deeper nonverbal, intuitive, or holistic ideas.

Sometimes the distinction is clear. Doing arithmetic, for example, or formally verifying a syllogism, are verbal/logical activities requiring little or no nonverbal or intuitive thought. The use of analogy or simile to capture or convey a complex or subtle idea, on the other hand, clearly involves a triggering of nonverbal processes, even though the communication itself may be purely verbal. In most situations, however, the distinction may be difficult to make, at least on the basis of the externally observable product of the verbal thought process. This is because careful nonverbal thought often produces a product that can be understood in purely verbal/logical terms. A mathematical proof, for example, can be verified by logic alone, even though an intuitive leap may have been necessary to produce it.

If the verbal product truly stands alone on logical grounds, as does, say, a valid mathematical proof, then the question of the role of nonverbal thought in producing it may be neglected. If, on the

other hand, it requires acceptance of some intuitive or unarticulated judgments, the validity of the overall product may depend critically on the quality of nonverbal thought that guided the production of that verbal product. This dependence is not necessarily diminished in any way by internal consistency within the product itself.

These two roles for verbal thought closely parallel the surrogate/perspective distinction made earlier for models. At the purely verbal level, words serve as surrogates for the ideas they express. Statements can be interpreted literally, and the ideas understood, and judged, in terms of those literal interpretations. If the language is being used to express a deeper nonverbal understanding, however, this may not be the case. The words may serve as perspectives on the deeper, more holistic ideas. The full meaning of the ideas, then, may not be adequately conveyed in the words alone. When words are used in this way, communication depends on the existence of an adequate, shared understanding of the context being addressed at the nonverbal level. Such understanding is often appealed to in policy analyses as a way around obvious literal weaknesses in the models used.

It is frequently the case in policy analysis that the substantive issue of real interest is extremely squishy and ill-defined, without sufficient unambiguously relevant natural structure to dictate a unique analytically tractable formulation. In such cases, the nonverbal intuitive thought involved in deciding how to analyze the problem, what problem elements to consider relevant and why, and in interpreting the logical results of the analysis in the light of those choices may play a far more important role in determining the knowledge gained from the analysis than do the strictly logical and computational portions of the analysis. No matter how faultless the logic, if the initial premises are faulty or ill-chosen, the conclusions will be unsound.

One of the major differences between mathematics and the physical sciences on the one hand, and policy analysis and other forms of operational inquiry into squishy questions on the other, is that the former are able to impose a combination of validity criteria and constraints on the set of problems considered that are sufficient to decouple

questions of validity of knowledge produced from the role of the intuitive nonverbal thought process in producing that knowledge, while the latter are not. Attempts are sometimes made to transfer methods, validity tests, etc., from the former to the latter that fail to take that difference into account. The role of nonverbal thought in understanding squishy problems is sufficiently critical, however, that it must be considered in addressing questions of how to deal with those problems.

The Cognitive Spectrum

We can think of human thought in general, and problem solving in particular, as taking place along a cognitive spectrum something like that shown in Fig. 11. The lower half of the spectrum represents internal human thought--that taking place within the mind. Nonverbal thought is at the lower end, with verbal above. Nonverbal thought consists of a number of different and deepening layers, but no attempt will be made to distinguish between them here. The shift from nonverbal into verbal is fuzzy and somewhat indistinct, as thoughts are imprecisely "put into words." The verbal lies between the nonverbal and the internal/external boundary.

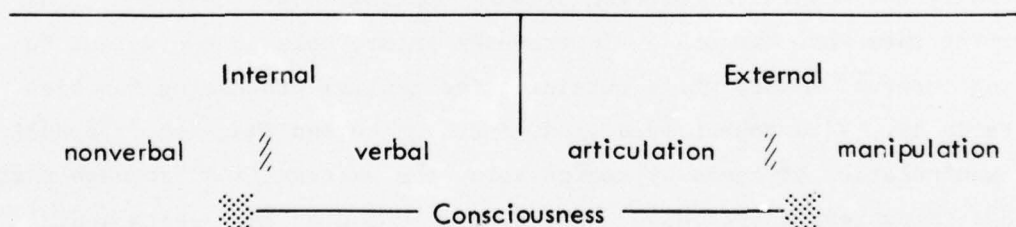


Fig. 11— The cognitive spectrum

Human thinking abilities, i.e., the ability to process thought, are not limited to what we can do within our heads. We have developed abilities to use external aids to extend thinking, and in so doing have made these aids, in essence, a part of our thought process. This

is indicated by the external portion of the cognitive spectrum in Fig. 11. The simplest use of external aids is the articulation of thoughts in order to keep track of them for bookkeeping purposes. When I use pencil and paper to write down numbers in order to add them up, for example, the pencil and paper serve as an external extension of my verbal thought process. They allow me to keep track of verbal thoughts (numbers) more permanently and reliably than I can in my internal memory. The processing of those numbers (forming the sum), however, is still one that takes place within my head. The same type of process is going on as I attempt to articulate more complicated thoughts in order to organize and keep track of them. Writing this report serves that function (in addition to the communicative function of making the thoughts available to others).

External articulation of ideas is also a primary mechanism of communication in forms of inquiry, such as scientific research or policy analysis, where the knowledge is produced for and used by others. The utility of the knowledge may depend heavily on how well both the knowledge and supporting grounds are articulated.

I can write things down (in words or numbers, or even nonverbal ideas in the form of sketches, etc.) as a way of expressing, recording, or keeping track of the ideas represented, while still retaining internally the cognitive function of manipulating those ideas and transforming them into new ones. In computer terms, this is equivalent to using external memory while retaining the central processing function internally. With appropriate aids, I can go beyond that--to externalize my manipulation of ideas by manipulating the externalized language that I use to express those ideas. The digital computer represents a pinnacle of this form of external manipulation. Simpler types of external manipulation, however, have been around for a long time in the form of adding machines, mechanical card weaving looms, etc.

Consciousness, or conscious awareness of our cognitive process, spans only a portion of this spectrum. The span of consciousness varies from person to person, running from somewhere in the upper portion of the nonverbal part of the spectrum through the low end of a manipulatory portion of the spectrum in most of us. In other words, we are conscious

AD-A031 629

RAND CORP SANTA MONICA CALIF

F/G 5/1

A CRITICAL ASSESSMENT OF QUANTITATIVE METHODOLOGY AS A POLICY A--ETC(U)

AUG 74 R E STRAUCH

P-5282

NL

UNCLASSIFIED

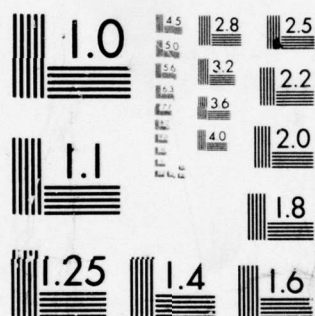
2 OF 2

AD
A031 629



END

DATE
FILMED
12-76



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

of all of our verbal thought, and of some portion of our nonverbal thought. We can also apply our consciousness to external expressions of our thoughts, to our language, drawings, etc., and can consciously follow simple manipulations of external thoughts such as those performed on an adding machine or abacus. We are unconscious of most of our nonverbal thought processing on the one hand, and are equally unconscious of most of what goes on inside of our computers on the other. Most creative thinking, and our "real understanding" of those things we "really understand," probably resides at or perhaps slightly below the lower limits of consciousness, in the upper reaches of our nonverbal thought.

We may accumulate and process information at various points on this spectrum and may transfer information from place to place. Information may be accumulated from internal as well as external and nonverbal as well as verbal sources. Internal sources include our memory and our existing internal cognitive models. External verbal sources of information include books, talking to people, computer printouts, etc. In addition, we obtain external information nonverbally through sense, experience, and impression, etc. We move this information up and down the spectrum. We put our nonverbal thoughts into words internally, then articulate them externally, then perhaps combine them in a computer with other external data and manipulate them externally. Going back down the spectrum, we read what is written, and we internalize nonverbally that which we read or think about verbally.

Internal processing may be purely nonverbal, a combination of verbal and nonverbal, or purely verbal. Purely nonverbal thinking--the processing of nonverbal thoughts, ideas, images, etc., without assigning verbal symbols to them--is something everyone does subconsciously. Some people, but not all, do some of it consciously. Purely verbal thinking involves the manipulation of ideas by manipulating the verbal symbols representing those ideas (words or numbers). Doing arithmetic is a purely verbal thinking activity, for example, as is verifying a logical argument, syllogism, etc. Purely verbal thinking is almost entirely a conscious activity. Most thinking is neither purely verbal nor purely nonverbal but is a combination of the two, involving the

simultaneous manipulation of nonverbal ideas and verbal symbols for those ideas. We are conscious primarily of the verbal component, although there is also a significant nonverbal component.

External thinking--the external manipulation of ideas as in a computer--is a purely verbal activity. It is done according to rules and procedures for manipulating symbols that may be related to, but strictly speaking are independent of, the meaning of the symbols. Low-order manipulations, such as those using a slide rule or adding machine, say, may be within the span of consciousness of the person doing the manipulation. More complex manipulations, however, such as those involved in large-scale computing, are as far outside of the consciousness of the manipulator as are the deepest nonverbal processes of the subconscious.

Verbal thinking, internal or external, is a process of manipulating and transforming symbols into new symbols. As such, the validity of the product of verbal thinking depends on (1) the consistency of the meaning assigned to the symbols throughout the manipulation and (2) on the validity of subjecting that meaning to the particular manipulations made. The validity of arithmetic, for example, depends on the fact that the symbols we use for numbers and arithmetic operations have consistent, well-defined meanings and can be validly subjected to the transformations we make in doing arithmetic. We can add things in either order, for example, because of the commutative law of addition. Fallacious syllogisms, such as "St. Louis ballplayers are Cardinals, hence, are members of the hierarchy of the Roman Catholic Church," are produced by a violation of the first requirement--consistency of meaning. The meaning of the word "cardinal" shifts during the manipulation. An inference such as "Joe is Tom's friend, Tom is Bill's friend, therefore, Joe is Bill's friend," may be wrong because of a failure of the second requirement--validity of the manipulation. Friendship is not necessarily transitive.

The risk that violation of these requirements will produce invalid conclusions is always present in verbal thinking, whether that thinking is internal or external. When it is internal, our consciousness of the thought process occurring and our nonverbal awareness of the meanings

of the terms being manipulated can serve as a check on the process and as a means of guarding against these errors. When the manipulation is external, however, outside the range of our consciousness as in a computer core, these same checks do not exist. The risk of fallacious inference in external thinking is increased by their absence, because we have lost conscious control of what's going on.

The ordering represented in the spectrum, at least up to the point of external articulation, seems reflective of a number of natural orderings. In moving up the spectrum, we find an increase in explicitness, precision, repeatability, use of logic in relating grounds to conclusions, etc. The ordering also seems to represent a natural evolution in human thinking at several levels. The development of an individual's thinking process, for example, seems to follow this spectrum. Young children initially think nonverbally, perhaps with very little conscious thought as to what they are doing. As they grow older, they gain the ability to structure their thoughts, knowledge, etc., using language. They learn to express their thoughts externally with increasing precision, and eventually, they learn to use external manipulatory devices.

In mature individuals, we place a qualitative valuation on knowledge according to the place on the spectrum its grounds lie. We give least value to "gut feelings" arrived at subconsciously and nonverbally, with no conscious understanding of their grounds. Moving up the spectrum, we place more value on well thought through opinion, and even more on well thought through opinion with externally articulated grounds. This same hierarchy is reflected in our view of cultural progress, as measured as the sources of authoritative knowledge in a culture and of the grounds for that knowledge. In primitive cultures, the authority may be a shaman, witch doctor, oracle, etc., who forms conclusions while in a trance or under the influence of drugs. He may describe this process as one of communion with the gods, spirits, or whatever. In terms of the cognitive model outlined here, he is letting his conclusion form subconsciously at the nonverbal level and articulating only the conclusion and none of the grounds. As society matures, we see an increasing evolution toward explicit verbal articulation of the grounds on which authoritative knowledge rests, culminating in the most

authoritative source of knowledge in contemporary society--science. In science, we ask that both the knowledge and the grounds be fully and explicitly articulated, open to review and criticism by others. We ask, in effect, that the grounds for knowledge be placed near the high end of the external articulation portion of the spectrum.

SHAMANS AND SYSTEMS ANALYSTS

The gradual emergence of science and rationalism as sources of authoritative knowledge is the result of an increasing requirement that the grounds for knowledge be made explicit--be made available to scrutiny and understanding--and that the individual proposing a solution to a problem be able to clearly demonstrate how he got that solution and why others should accept it. Until the advent of the digital computer, this requirement was met by pushing the grounds for knowledge as far up the spectrum of explicitness and precision as possible. Until the computer, the upper end of the spectrum was still within the span of human consciousness. The computer, however, has given us an ability to manipulate ideas far beyond our ability to keep track of and understand what we are manipulating.

Even the author (if there is a single identifiable author) of a large computer program seldom understands fully what it does. Additional users are likely to understand it even less. The individual pieces of a program can be understood and can be verified, but the way they fit together, and the way the program works as a whole, is frequently beyond the grasp of any of the users.* In effect, the mechanism producing the knowledge is again one that must be taken on faith, as it is in the case of the shaman or the oracle. It is faith in a different mechanism--in machines and equations rather than gods or spirits. It is a faith in the belief that the validity of a conclusion can be insured by insuring the validity of the assumptions one by one--each of the little pieces. It is the belief that if each

* For an excellent discussion of this point, see "On the Impact of the Computer on Society," by Joseph Weizenbaum, *Science*, Vol. 176, May 12, 1972, pp. 609-614.

individual assumption is valid, then the way they fit together and the conclusion they produce will be valid also. This faith may be born out if the model used is a sufficiently close fit to the problem being addressed, i.e., is an adequate surrogate for that problem. Witness, for example, our ability to make the trajectory calculations necessary for our moon shots, or to run large petroleum refineries with a minimum of human intervention. In the case of many of the squishy problems encountered in policy analysis, however, our knowledge is insufficient to allow the construction of adequate computer-based surrogates. When we look for answers to complex social, political, or behavioral problems in computer based models, therefore, we may be returning full circle--back to the shaman and the oracle, asking a magical mechanism beyond the range of human consciousness and understanding to provide us our answers.

I have in the past occasionally made a semi-facetious comparison between the shaman and the systems analyst--between method-oriented quantitative analysis and the reading of entrails. As I thought more about it, however, I decided that the analogy is less facetious than I first believed. Key roles in the analogy are played by the computer, which serves as the contemporary substitute for the chicken, and by the results it produces, which are the contemporary entrails. As the shaman sliced open his chicken and learned from the viscera that fell to the ground, so the systems analyst with the large complex computer model examines the results that the computer spews out and prognosticates wonderful things. In his own way, each appeals to a power beyond the ken of ordinary mortals and outside the span of human consciousness--a power that must be accepted because of the homage paid to it by their respective societies.

I am not suggesting that computers have no role to play in extending the bounds of human knowledge. They certainly do. Their appropriate role, however, is as an extension of, rather than replacement for, human judgment. When they are used to replace judgment, to manipulate symbols according to abstract rules in problems that we understand insufficiently to prescribe abstract rules for, they become a threat to understanding rather than an aid. They do this in two ways.

First, they fail to provide the insight that they promise and that they superficially appear to deliver. They cannot provide "optimal strategic force postures" or "optimal energy policy" or even "a best estimate of the elasticities among inputs" in complex political situations, simply because man does not now understand those problems sufficiently to give those terms precise quantitative meaning. But in spite of their inability to deliver, however, they promise. If that promise is accepted for what it claims to be, the ability of the acceptor to truly understand the problem is diminished accordingly. When this happens, the computer has not only served as an inadequate instrument for providing insight, but has done positive harm by blinding its constituency to the knowledge that may be available through more careful thought (with or without computer assistance).

Consider, for example, the use of computer models of NATO/Warsaw Pact conflict in analyzing NATO military requirements and the security problems NATO poses for the United States. What might we reasonably expect to learn from such models, and what, in fact, do we learn or think we learn? Are these in balance, and if not, what imbalances may exist and why?

There are many computer models of NATO/Warsaw Pact conflict in use. They range, in level of detail considered, from very aggregate models using gross force balances to highly detailed models reproducing combat interaction down to a company, platoon, or even squad level. Some are simulations, some expected value models, and some deterministic. Different models may emphasize different factors, depending on their design objective. Some, for example, might emphasize the role of armor, others the role of air, others the role of logistics, etc.

In spite of these superficial differences, all models of this type embody essentially the same approach to the study of NATO war. The model is a description of war as a mathematical system. This system specifies the inputs (force levels, accuracy, yields, damage expectancies, etc.) required as initial conditions, the outputs by which the outcome of the conflict is defined (area taken, casualties produced, tons of ammunition expended, etc.) and the relationship between the two that determines how inputs produce outputs. Runs of the

model use the inputs and relationships to produce the outputs--the results of the conflict process described by the mathematical system. Any interpretation of these results as conclusions about a "real" conflict process embodies assumptions about the relationship between the mathematical system described by the model and that "real" process. The validity of any conclusions will depend strongly on that model/process relationship.

The output produced by such models is frequently in terms of war outcome, or some fairly direct proxy for war outcome. One such proxy is the movement of the forward edge of the battle area (FEBA). Since computational outputs such as FEBA movement appear to predict war outcome as a function of the input variables, one natural question to ask about such models seems to be:

Can computer models of NATO/Warsaw Pact conflict reasonably be expected to produce useful and valid predictions of the outcome of such a conflict?

By *valid*, I simply mean predictions such that it seems reasonable to believe, with some degree of confidence, that a real conflict with initial conditions as specified in the model would produce an outcome approximating that produced by the model. By *useful*, I mean that the model is able to make such valid predictions in circumstances, sets of conditions, etc., in which equally valid predictions are not readily available by other, simpler means.

Potential NATO/Warsaw Pact force balances fall into three categories, as shown in Fig. 12. At one extreme, the balance may so overwhelmingly favor NATO as to give NATO a clear dominance in any NATO/Warsaw Pact conflict. At the other, the balance may so strongly favor the Warsaw Pact as to give them clear dominance. Between these two extremes lie a range of force balances having the property that neither side is so evidently superior to the other as to have a clear dominance. The outcomes from conflicts in this middle range are not clearly determined. The two extremes of clear NATO or Warsaw Pact dominance could be ascertainable without resort to complex computer models. For force balances in these categories, the models may give

valid outcome predictions, but the predictions are not particularly useful. They can be readily ascertained by simpler means. These don't represent the very interesting cases anyway, since neither side is likely to allow the balance to drift to one of these extreme states.

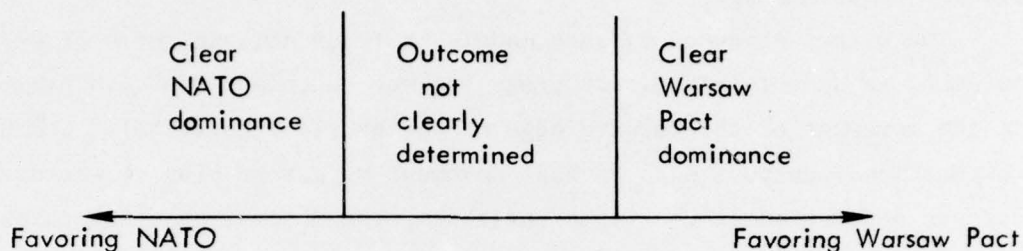


Fig. 12 — NATO-Warsaw Pact balances

But what about the range in the middle, where the NATO/Warsaw Pact balance is such that the outcome is not clearly determined? Here, if anywhere, is where complex modeling should pay off. The known characteristics of the forces on both sides can be fed into the model (at whatever level of detail the model requires) and the model will produce a predicted war outcome. But how valid is it likely to be? If we are to consider the prediction valid, a minimum we might ask is that it be sensitive to the effect of variables that have historically proven themselves to be important in determining war outcomes in conflicts involving ambiguous force balances. If the model is not sensitive to such variables, then we should seriously question its predictive validity. One variable that has proven itself repeatedly to be of major importance is *good generalship*--the ability to decide how to use the forces as the battle progresses in ways that take advantage of the enemy's weak points and mistakes. One seldom sees, however, a computer campaign analysis model that even recognizes the existence of, let alone is sensitive to, good generalship. This being the case, it seems reasonable to question the validity of predictions produced by such models.

The answer to our question, then, would seem to be "No!" In the case of clear dominance by one side or the other, computer models can probably produce valid predictions but such predictions are not particularly useful. Where the outcome is not clearly determined, computer models universally ignore one of the major factors in determining war outcome--good generalship. Hence, the predictions they make are of questionable validity. The models, therefore, promise a product (valid and useful war outcome predictions) that they are incapable of producing.

It might be argued that this is a straw man--that designers and users of computer campaign models are the first to admit to the inability of their models to predict war outcomes, but that such models are still useful for making relative comparisons between alternative weapons or weapon systems, for understanding the effects of the introduction of new capabilities, etc. As a counter argument, I would reply that the prominence usually given to war outcome predictions (caveats in the fine print notwithstanding) is such as to make their validity a relevant subject in questioning and evaluating the utility of the models. If the model designer or user really believes in his caveats--that his results should not be interpreted as predictive of what would actually happen--one might think he would give a little less prominence to outcome *per se*, and a little more to whatever it is he believes his analysis really shows. If the results are displayed as a series of battle maps showing Warsaw Pact forces rapidly advancing to the Elbe, for example, a listener can hardly be faulted if his primary impression after the briefing is one of "how fast they took Hamburg." The counter-counter argument that these are simply matters of interpretation or presentation is insufficient, since presentation and interpretation are what determine what the analysis has to say about the substantive problem.

But what of the arguments that the models are useful for relative comparison or for evaluating the effect of changes in capability? Are those arguments sound or unsound? That depends on how the models and the results they produce are used, on the relative emphasis given to model results *per se*, and to their careful interpretation. Direct

translation of model results into substantive relative comparison--"the outcome is unchanged if we add one tactical air wing and take away two armored divisions. Therefore, we can trade air wings for armored divisions on a one for two basis," say--is as likely to be totally specious as is absolute war outcome prediction. More subtle comparisons, such as might be obtained by understanding why the outcome remains unchanged, and relating that reason more directly to "real" combat, may be valid. This requires more than the model alone provides, however. It requires careful and considered explicit interpretation of the model as a perspective on (rather than a surrogate for) the complex and squishy process of land combat. It requires that the results of the model--produced by external manipulation of language beyond the limits of human consciousness--be brought back into consciousness in the human mind, and perhaps, down into the deeper non-verbal parts of the mind for interpretation as substantive conclusions about a real process. If the model is an inadequate surrogate for the problem, then the important conclusions are those that the *analyst draws* from the model--not those that the model produces. The difference is a subtle but important one.

The long-run danger in overemphasizing models that neglect the value of good generalship, and in undervaluing the need for careful human judgment in evaluating or interpreting the results they produce, is far more serious than the risk of a few erroneous study conclusions. The long-run danger of such an approach lies in the fact that we may not have the good generals when we need them. Understanding war, and understanding computer models of war, may be two very different skills. Contemporary experience revolves increasingly around the latter. If we fail to distinguish carefully enough between the two, we may create and reinforce an environment in which the incentive and reward mechanisms increasingly favor those who understand computer models of war and who derive their understanding of war from such models, and increasingly exclude those who understand conflict in other ways. The long-run effect of this, over, say, a generation, could be to produce a military establishment eminently qualified to fight the wars embodied in their computer planning models, and largely incapable of adapting to anything else.

Because of the lack of a continuing stream of "real" experience, this sort of risk is perhaps greater in defense planning than it is in other segments of government or the larger society. It is by no means absent elsewhere, however. It is, in fact, inherent in any approach to knowledge that places the grounds for knowledge outside the bounds of human understanding--that accepts conclusions on faith because of the might and magic of the power that produces them. This is as true when that power is the "computer" as when it is the "Genie in the Eight Ball" or any other deity.

A MODEL OF JUDGMENTAL ANALYSIS

The model of judgmental analysis and problem solving suggested by the preceding considerations is sketched in Fig. 13. The external

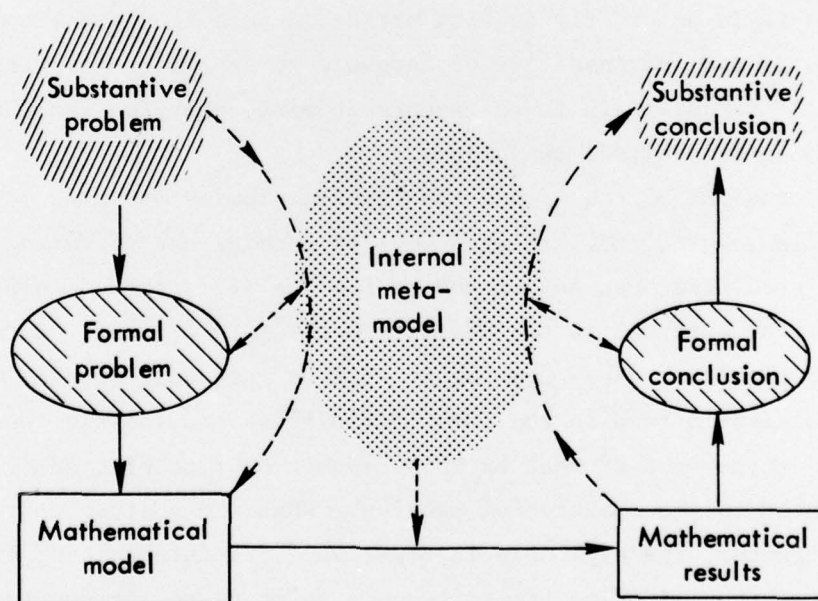


Fig. 13 — A model of judgmental analysis

portions of the model, i.e., the elements of analysis accessible to an external observer, are the same as in Fig. 3, and are shown in solid lines. The internal components--the analyst's internal meta-model and the connection between the meta-model and the external structure--are shown in dotted lines.

This model assumes the existence of a single analyst responsible for the overall analysis--beginning with the substantive problem and ending with the substantive conclusion. The complications introduced by team analysis, or analysis split into different levels performed by different people, are severe, but will not be addressed here. Also, the present discussion will treat the process of analysis as a sequential process in which the elements of Fig. 13 are constructed one at a time in orderly fashion. The reader should remember that the process of analysis is seldom like this. Rather, it is a complex and disorderly process while it is occurring--which can be neatened up and put into the orderly sequence discussed here only at the end. The need to describe it in an orderly fashion arises in part from the linear nature of language and the inability of language to deal directly with complex and subtle ideas. With these caveats in mind, then, we can think of the process of analysis as follows.

The analyst begins with a substantive problem and a collection of information, data, etc., related to that problem and of methods, techniques, etc., that can be brought to bear on it. Some of this information will be external to the analyst in the form of facts, data, other people's opinions, etc. Some of it will be internal--already stored in the analyst's head in the form of a diffuse and loosely constructed internal *meta-model* of what he knows about the type of situation represented in the substantive problem. What the analyst wants to do is to structure the available information (both internal and external) in a way that will allow him to reach a solution to, or conclusion about, that problem.

He begins by absorbing at least some of the information defining the substantive problem and combining this with knowledge he already possesses about that problem and about available methods, techniques, etc. that he may bring to bear on its solution. He formulates models

of the problem on at least two distinct levels--the formal and the mathematical. He may or may not explicitly distinguish between the two. His mathematical model is an abstract mathematical structure whose logical implications he will explore as part of his mathematical analysis. The formal model provides him with a structuring of the substantive problem and its elements which serves to link his mathematical model with the more diffuse substantive problem. The mathematical model provides an external logical aid for following and developing the logical implications of a specific set of premises.

These are the activities referred to earlier as *formulation*, and reflected by the dotted lines on the left side of Fig. 13. (Ignore the solid lines for the time being. We will get to them later.) Formulation is essentially a judgmental activity taking place inside the analyst's head. I will not attempt to specify logically how it is done, but I do want to include the fact that it is done explicitly in my model of analysis.

The model is analyzed mathematically, producing the mathematical results. The analyst takes these results back into his head--into his internal meta-model, if you will--to combine them with the additional knowledge contained there about the substantive problem, which was not included in the model. In combination with this knowledge, the results are given interpretations at two levels--the formal and the substantive, with the formal again serving as an intermediate link between the mathematical results and their interpretation as a substantive conclusion.

Throughout this process, internal thought processing takes place at both the verbal and nonverbal levels. Much important processing occurs at the nonverbal level (possibly subconsciously) with verbal thought serving to express that nonverbal thought, and to provide an interface between it and the external structure. The intermediate verbalizations of the formal problem and the formal conclusion serve to provide external structures that will hold still to be consciously examined and checked against other available information, etc., better than the analyst's internal model will.

The linkage in Fig. 13 between model and results is solid, indicating that, in principle, at least, this linkage is external or

could be made so. Depending on the model and the particular results derived from it, this linkage may consist of a logical chain of proof connecting premises with results, of numerical calculations within the framework of the model, or both. If numerical calculations are involved, these may result from the analyst's internal thought processes (e.g., on the back of an envelope), or from the use of an external mechanism such as a computer. In the former case, the calculations will necessarily be within the span of consciousness of the analyst making the calculations, while in the latter, they may be outside that span. However produced, the results are tautologies within the structure of the model. Hence, in some very real sense, the mathematical analysis is the intellectually most trivial portion of the analysis.

While the link connecting the model and results is necessarily a verbal/logical one, the nonverbal/intuitive thought process plays an important role in the construction of that link. If the link is a logical chain of proof connecting premises and results, then the analyst's use of his intuition and internal knowledge about that model, similar models, techniques, proofs, etc., will likely play an important role in directing him to the results obtained. If he's lazy, he may trust his intuition too far and fail to complete the external logical proof. If he does so, he runs the risk of erroneous results.

The other way that the nonverbal/intuitive process guides the mathematical analysis is in determining what results are worth having, in terms of the substantive problem. In this role, the analyst uses not only his internal knowledge of the model and of analytical techniques, but also his understanding of the substantive problem and of the problem/model relationship. Unless the model is a very good fit to the problem, some portions of the model are likely to be attributable to simplifying assumptions chosen to make the problem tractable, rather than to reflect the real world structure of the problem. Results driven by these parts of the model are likely to be of little substantive value even though they are valid results in the mathematical sense.

The analyst who fails to exercise his subjective judgment about which results are likely to be meaningful, given the existing problem/model relationship, runs the risk of wasting his time in producing

useless results. He is also likely to be guilty of sloppy or erroneous interpretations, since a lack of judgment in mathematical analysis and in interpretation often go together. The risk of error of this type is particularly high in complex computer calculations, since the calculations themselves are likely to be beyond the span of consciousness of the analyst, thus, difficult to keep track of and understand at the intuitive level.

The surrogate/perspective distinction comes into play here. If the analyst is thinking of the model as a surrogate, then he may feel that any results that are mathematically valid within the context of the model also have substantive meaning. If, on the other hand, he is thinking of the model as a perspective on the substantive problem he is likely to find some results within the model context of more substantive value than others, even though they may all be of equal logical validity.

Having produced the results, the analyst takes them back within his internal model for interpretation (shown by the dashed arrows on the right side of Fig. 13). Like formulation, interpretation is a judgmental process. While I will not attempt to prescribe logically how it is done, it is an explicit and important part of this model of judgmental analysis. It involves the integration of the results with additional internal knowledge possessed by the analyst about the model/problem fit. This includes his knowledge and judgments about aspects of the substantive problem that are not reflected in the model and about aspects of the model structure and behavior that are not representative of the problem. It produces acceptance of some results as valid formal or substantive conclusions, qualification of others to account for problem/model differences, and rejection of others as the product of discrepancies between the model and problem.

The surrogate/perspective distinction plays an important role in interpretation. If he is using the model as a surrogate, the analyst is likely to make fairly direct translations from results to conclusions--to interpret his results directly as formal and substantive conclusions. If, on the other hand, he is using a model as a perspective, he is more likely to qualify his interpretations and to

distinguish between what the results say about the model per se and what that implies about the substantive problem.

Documentation and Communication

Analysis is more than just reaching a conclusion by the process just described. It also involves the documentation of that conclusion--the production of a permanent external version of the conclusion and grounds. This documented version serves the following three functions:

1. It allows the analyst to formalize and verify what he has done. It gives him a permanent structure that will sit still while he looks at it and makes sure he is satisfied with it, in a way that the internal structure in his head will not.
2. It provides a basis for the communication of his conclusions and of the grounds for those conclusions to others--customers, clients, peers, etc.
3. It provides a basis for judgment by others about the conclusions he has reached and the validity of those conclusions.

In reaching his conclusions, the analyst may have worked primarily in his head--within his internal meta-model--articulating externally only those parts of the structure shown in Fig. 13 that he needed to articulate in order to keep track of what he was doing. Documentation of his conclusions and grounds involves the articulation of a great deal more of that structure--of as much as possible of the portions of Fig. 13 shown in solid lines.

The process of documentation produces an external, verbal description of this structure and of the solid lines connecting its elements. It is this structure from which the reader, user, customer, etc., must determine what the analyst has done and evaluate its validity. The question of validity, per se, is one that depends on the problem, the conclusions, and the grounds for those conclusions, in a highly complex and problem dependent way. It is not one about which we can say much at the broad brush level of generality addressed here, where we are concerned primarily with the structural characteristics of analysis.

There are two structural characteristics of analysis and its documentation, however, that relate to validity (or at least to the problem of judging validity) which can be examined at the level of generality considered here. These are the *separation* of the conclusions from the analyst and the *comprehensiveness* of documentation.

Separation from the analyst is a property of the conclusion and grounds. It refers to the extent to which the conclusion is grounded in logic and objective fact, as distinct from the subjective judgment of the analyst--the extent to which the dotted lines in Fig. 13 can be severed and the conclusions and grounds made independent of the analyst who produced them. Following the usage already established in this report, a conclusion that can be completely separated from the analyst who produced it will be referred to as *free-standing*, while one that cannot will be referred to as *judgmental*. In terms of the cognitive spectrum shown in Fig. 11, a free-standing conclusion is one whose grounds can be moved up the spectrum to the point of external articulation, while a judgmental conclusion is one that remains grounded, at least partially, in the internal and perhaps nonverbal part of the spectrum.

Comprehensiveness is a property of the documentation. It refers to the extent to which the grounds for the conclusion are fully articulated rather than being left to the reader to fill in or infer. It refers not so much to the level of detail--more can always be said about a complex problem--as to how completely or incompletely all the important things are said and how much work the reader still has to do in order to understand them fully.

As in the case of the two meanings of objectivity discussed earlier, we sometimes think of these as characteristics that occur together--of comprehensively documented free-standing conclusions as opposed to sparsely documented judgmental conclusions. A fully documented, rigorous, scientific proof might be a prototype of the former, while assertions by an oracle or mystic about knowledge received while in a trance or under the influence of drugs would be an extreme prototype of the latter.

But like the meanings of objectivity, they need not occur together. Theoretical mathematics is perhaps the most completely free-standing

area of knowledge known to man, but anyone who studied it has often come across the phrase "with the details left to the reader," and knows how incomplete a proof that phrase can sometimes cover. Few of us require (or would be likely to tolerate) exposure to full documentation of most of what we accept as scientific fact. At the same time, numerous examples exist of comprehensively documented subjective analysis, in which the analyst has articulated his conclusions and grounds well, even though they remain strongly tied to his judgment.

There is a difference between judgmental and free-standing conclusions, perhaps, in the potential comprehensiveness with which the grounds could be articulated. In principle at least, the grounds for a free-standing conclusion may be articulated completely--built up from fundamental axioms and reproducible observations of nature. This level of comprehensiveness is not achievable, even in principle, with respect to judgmental conclusions. These remain ultimately grounded in the internal meta-model of the analyst who produced them. They can be fully appreciated and reproduced by a reader only to the extent that he shares a compatible meta-model.

A related difference between free-standing and judgmental conclusions is in how confident we feel with incomplete grounds. If we believe a conclusion to be free-standing, then we believe it is grounded in fact and logic and hence is "objective" in the sense of describing things as they are. We may, thus, happily accept the conclusion without fully understanding the grounds, confident that examination of those grounds in greater depth would only affirm our belief in the conclusion. This is not the case, however, with the judgmental conclusion. Even if it's our own and has been very carefully thought through, we can never be sure that more analysis and thought wouldn't change it. How much less confidence then must we have in someone else's incompletely documented judgmental conclusion?

For this reason, then, it seems natural to prefer free-standing knowledge to judgmental knowledge whenever possible. This preference has been strengthened by the advantages conferred on contemporary society by the application of free-standing knowledge in the form of science and technology. It is this preference, bolstered by the success

of science and technology, that has resulted in the shift from sage to scientist as the respected giver of authoritative knowledge, and from wisdom to technical expertise as the quality to be sought after and cultivated.

The principle that free-standing knowledge may be accepted without comprehensive documentation of its grounds, because its free-standing nature insures its validity, is a reasonable one in principal. In practice, however, it contains its own version of "Catch 22"--without fully examining the grounds, how can we be sure the conclusion is free-standing? One way is to accept the word of experts. That doesn't eliminate the logical problem, however, because it tends to make the conclusion itself judgmental, now grounded in the judgments of the experts involved. Carried to its logical extreme, this line of reasoning suggests that, in principle, there may be no such thing as free-standing knowledge. This conclusion finds support in Thomas Kuhn's study of scientific revolution.* Kuhn suggests that, to a far greater extent than most of us would like to admit, scientific knowledge is judgmental knowledge--grounded in a set of conventions that the scientific community of the day chooses to accept without question. He further suggests those conventions change with time, when scientific revolutions occur, and that after the revolution science rewrites scientific history to make it appear that the changes never took place--that the current conventions have always been true. It is not my intention here to question the foundations of physics. As we examine the structure of knowledge built on quicksand, however, it may be useful to remember that even that which appears grounded in solid rock may have shifting sand underneath.

At a more practical level, the desire to produce free-standing conclusions, or at least conclusions that appear to be free-standing, causes serious abuses of quantitative methodology in policy analysis. The methodology itself, grounded as it is in mathematical theory, is free-standing when applied to the well-defined problems treated by

*Thomas S. Kuhn, *The Structure of Scientific Revolutions*, University of Chicago Press, Chicago, 1962.

that theory. Its application to the more complex and squishy problems arising in policy analysis, however, is a judgmental process, and the conclusions reached remain essentially judgmental conclusions. If the analysis is done and documented in a way that stresses its mathematical component and suppresses the role of judgment in formulation and interpretation, however, this fact may also be suppressed. The conclusion may then superficially appear free-standing, or nearly so. Contemporary approaches to analysis and documentation (described earlier as method-oriented analysis and quantificationism) do this. They result in the production of superficially free-standing conclusions that are really quite heavily grounded in unarticulated and often unconsidered judgments about problem/model fit.

Perhaps the framework outlined here, which is still in a very primitive and preliminary stage of development, can usefully illuminate some of the issues involved and suggest directions in which their resolution might be sought. One direction that suggests itself is a shift in focus concerning what constitutes analysis, away from the narrow focus currently extant on the mathematical components of analysis, and broadened to include the overall process and the importance of the judgmental components of that process. This, in turn, would suggest approaches to analysis that gave clearer emphasis to its judgmental aspects and approaches to documentation that aim at more comprehensive documentation of the full set of grounds, rather than focusing on the narrow portion of the grounds contained in the mathematical analysis.

An image of the quantitative policy analyst as a giver of fact, derived from viewing analysis as an "objective" activity, was discussed earlier (page 34). In contrast to that image, the image suggested by this model of analysis is one of the analyst as a giver of insight and judgment. His role in performing analysis is to use the specialized tools at his disposal to understand and gain insight into his substantive problems, or at least those parts of the problems for which the tools do provide insights. He should then use that insight and understanding as a basis for the judgments that make up his substantive conclusions.

Documentation of the analysis should provide a comprehensive description of those conclusions and the grounds on which they are based. It should serve as a basis for the transfer of understanding of the conclusions and grounds from the analyst to the customer, and as a basis on which others can understand what judgments he reached and why, in order to criticize and validate or refute those judgments.

This is difficult to achieve even when analysis is performed by a single sensitive and intelligent human being. The difficulties increase considerably when the analysis is fragmented between numbers of people or different levels in an administrative hierarchy. One reason for this is that the more people who are involved the harder it is for anyone to have a full intuitive grasp of the overall activity. Yet this grasp, as we have seen, is essential to good interpretation. I have no pat answer to this problem, but I believe that more attention to the directions outlined in this report would be a step in the right direction. It seems clear, in any case, that there is a problem, and that pretending it is not there will not make it go away.

BIBLIOGRAPHY

Brunner, Jerome S., and Leo Postman, "On the Perception of Incongruity: A Paradigm," *Journal of Personality*, Vol. 18, 1949, pp. 207-223.

Kuhn, Thomas S., *The Structure of Scientific Revolutions*, University of Chicago Press, Chicago, 1962.

Lilly, John Cunningham, *The Mind of the Dolphin*, Doubleday & Co., Inc., Garden City, New York, 1967.

Ornstein, Robert E., *The Psychology of Consciousness*, Viking Press, New York, 1972.

Strauch, Ralph E., *The Operational Assessment of Risk: A Case Study of the Pueblo Mission*, The Rand Corporation, R-691-PR, March 1971.

-----, *The Fallacy of Causal Analysis*, The Rand Corporation, P-4618, April 1971.

-----, "Winners and Losers: A Conceptual Barrier in Our Strategic Thinking," *Air University Review*, Vol. XXIII, July-August 1972, pp. 33-44. Also available as P-4769, The Rand Corporation, June 1972.

Weizenbaum, Joseph, "On the Impact of the Computer on Society," *Science*, Vol. 176, May 12, 1972, pp. 609-614.